ECCC

Sum-of-Squares meets Nash: Optimal Lower Bounds for Finding *any* Equilibrium

Pravesh K. Kothari [*]   Ruta Mehta [†]

Abstract

Several works have shown unconditional hardness (via *integrality gaps*) of computing equilibria using strong hierarchies of convex relaxations . Such results however only apply to the problem of computing equilibria that optimize a certain objective function and not to the (arguably more fundamental) task of finding *any* equilibrium.

We present an algorithmic model based on the sum-of-squares (SoS) hierarchy that allows escaping this inherent limitation of integrality gaps. In this model, algorithms access the input game only through a relaxed solution to the natural SoS relaxation for computing equilibria. They can then adaptively construct a list of candidate solutions and invoke a *verification* oracle to check if any candidate on the list is a solution. This model captures most well-studied approximation algorithms such as those for Max-Cut, Sparsest Cut and Unique-Games.

The state-of-the-art algorithms for computing exact and approximate [LMM03] equilibria in two-player, n-strategy games are captured in this model and require that at least one of i) size ($\approx$ running time) of the SoS relaxation or ii) the size of the list of candidates, be at least $2^{\Omega(n)}$ and $n^{\Omega(\log(n))}$ respectively. Our main result shows a lower bound that matches these upper bound up to constant factors in the exponent. This can be interpreted as an unconditional confirmation, in our restricted algorithmic framework, of Rubinstein's recent *conditional* hardness [Rub16] for computing approximate equilibria.

Our proof strategy involves constructing a family of games that all share a common sum-of-squares solution but every (approximate) equilibrium of one game is far from every (approximate) equilibrium of any other game in the family. Along the way, we strengthen the unconditional lower bound against enumerative algorithms for finding approximate equilibria due to Daskalakis-Papadimitriou [DP09] and the classical hardness result for finding equilibria maximizing welfare due to Gilboa-Zemel [GZ89].

1 Introduction

How hard is it to compute equilibria in two player finite games?

This foundational question has been a driving force of research in algorithmic game theory for almost three decades. Beginning with the formalization of the complexity class PPAD, a systematic investigation [Pap94, AKV05, SvS06, DP09, DGP09, CDT09] of this question eventually led to the elegant result [DGP09, CDT09] establishing that computing equilibria (NE) in two-player games is PPAD-complete.

Algorithmic progress has thus relied on making additional assumptions on the structure of the games or relaxing the notion of equilibria [DP09, LMM03, AKV05, KT10, Bar15, TS07, AGMS11, BR17]. Perhaps the biggest success in this direction has been the celebrated quasi-polynomial time approximate scheme of Lipton, Markakis and Mehta [LMM03] for approximate NE. In a major recent breakthrough, Rubinstein [Rub16] showed that we cannot obtain any significant asymptotic improvement over this algorithm assuming a strong, new conjecture: solving the PPAD-complete "End-of-the-line" problem requires exponential time.

We thus obtain a fairly complete picture of complexity of NE in general games if we believe strong enough fine-grained conjectures about the complexity of PPAD. This work is focused on the central goal of obtaining reliable, independent evidence (given that a resolution currently appears out of reach!) for such fine-grained complexity assumptions.

Unconditional Hardness for NE**?.** A standard approach for investigating such questions attempts to build a *conditional* theory of hardness based on (more believable) conjectures. This approach has the attractive feature of understanding the limitations of any algorithm for computing NE. However, since it is known that NE cannot be NP-hard unless NP = co-NP [MP91], such a plan encounters serious obstacles.

The other influential approach builds *unconditional* hardness results for strong but restricted class of relevant algorithmic techniques. Such results have largely focused on algorithmic techniques based on systematic hierarchies of linear/semi-definite programs (LP/SDP), such as Sherali-Adams, Lovász-Schrjiver and Sum-of-Squares. The focus of this work is investigating the sum-of-squares SDP hierarchy for computing NE.

The sum-of-squares SDP hierarchy generalizes spectral methods and linear programming and is formally the strongest known hierarchy of general, convex-programming based algorithmic schemes. It captures the state-of-the-art algorithms for many fundamental algorithmic problems [ABS15, BRS11, GS11, ARV09] including the celebrated QPTAS for computing approximate equilibria in game [HNW16], breaks many known hard instances of basic problems in combinatorial optimization [BBH+12, OZ12], and has been remarkably successful in algorithm design for both worst-case and average case problems [BKS17, BKS14, MSS16, BKS15, BM16, KS18b, KS18a, BBKK18, HL18, KKM18].

Given the power of the technique, lower bounds against it can be credible indicators of computational hardness and have been successfully used as such especially in areas where standard conditional hardness results are infeasible. For example, some of the strongest known evidence of hardness of many fundamental *average-case* problems such as Planted Clique, Refuting Random Constraint Satisfaction Problems and Maximizing Random Polynomials (or, "Tensor Principal Component Analysis") comes from strong lower bounds [BHK+16, AOW15, KMOW17, RRS16, HSS15] against the sum of squares algorithm.

However, a priori, this approach also suffers from a similar foundational problem as the first one. Lower bounds for convex relaxations are usually formalized as *integrality gaps*. Integrality gaps show impossibility of approximating a certain objective function on the underlying solution space with the given convex relaxation. However, NE is a search problem without any pre-defined objective function. Restricting to the natural decision version trivializes the problem as every two player finite game has an equilibrium. Thus, integrality gaps do not provide means of establishing hardness for the problem.

Prior works [CS08, HK11, BKW15, DFS16] dealt with this issue by forcing some natural objective function on the space of equilibria and establishing the hardness of finding an equilibrium that optimizes the chosen objective. For appropriate choices of such objectives, finding (or even approximating) the best equilibria is known to be NP-hard [GZ89, CS08]. However, such a result doesn't allow distinguishing whether the hardness is of finding *any* equilibrium or one that optimizes the chosen objective function.

The main contribution of this work is a new approach to circumvent such issues and establish unconditional hardness for finding *any* equilibrium using the sum-of-squares method by relying on *rounding gaps* instead of *integrality gaps*.

Hardness via Rounding Gaps. The conceptual idea underlying our framework is quite simple. A convex relaxation, such as the SoS SDP, returns a certain relaxed solution for the problem of our interest. A relaxed solution will not generally be an actual solution so we thus must use an additional, second step, usually known as *rounding*, that transforms a relaxed solution into an actual equilibrium. Given that any algorithm for NE based on such a convex relaxation must go through the additional rounding step, we will show hardness results for both the steps above combined.

There's an important issue, however. At the outset, the rounding algorithm might just ignore the convex relaxation's solution and just find an equilibrium for the underlying game from scratch. Without settling the complexity of PPAD, thus, we cannot hope to prove a lower bound against such "rounding" algorithms. Thus, we must restrict our rounding to procedures that really "use" the solution to the convex relaxation and do not "cheat" by ignoring it.

Oblivious Rounding and Verification Oracles. An *oblivious* rounding algorithm takes input the relaxed solution (generated by the convex relaxation) and outputs a true solution with the crucial restriction that the output be a function of *only* the relaxed solution. In particular for us, an oblivious rounding algorithm for computing NE accesses the underlying payoff matrices of the input game only indirectly via the solution to the SoS SDP.

This class of roundings was first formalized in a work of Feige, Feldman and Talgam-Cohen [FFT16] in the context of certain mechanism design problems (see a longer discussion in Section 2.3) . While such rounding algorithms might appear restrictive at the outset, a simple inspection reveals that, in hindsight, some of the most powerful SDP rounding algorithms in the literature are in fact oblivious! This includes, for instance, the famous algorithms for Max-Cut [GW01], Sparsest Cut [ARV09] and unique games/small-set-expansion [BRS11]. Finally, a recent work of Harrow-Natarajan and Wu [HNW16] shows that the guarantees of the Lipton-Markakis-Mehta [LMM03] algorithm for computing approximate equilibria can be matched by an *oblivious* rounding based algorithm.

While oblivious rounding captures many famous rounding algorithms, there are a few notable exceptions that do not fit into this framework. This includes, for e.g., the rounding method for

arbitrary Constraint Satisfaction Problems [RS09] and the recent works on polynomial optimization and its variants [BKS14]. Such works construct a list of candidate solutions based on the solution to the relaxation and show that one of the candidates must in fact be a true solution. Indeed, such an idea in particular captures the spectacularly successful "enumerative techniques" in computing equilibria for various classes of games: cleverly construct a small search space that contains a solution and then brute-force search over it to find the solution, e.g., [LMM03, KT10, DP09, ALSV13]. Perhaps the most famous example of the usage of this idea is in fact the algorithm in [LMM03] discussed above.

We thus allow our model to (possibly adaptively) generate a list of candidates and check if one of them is a solution. This gives us the final class of roundings we work with in this paper: *Oblivious Roundings with Verification Oracle* (OV). We measure the cost of the rounding in two parameters: the running time of the relaxation and the number of candidate solutions generated. Notably, we do not restrict the running time of the algorithm that constructs such a list from the solution to the relaxation - in that sense, our model is information theoretic. We also allow certain strengthening of the verification oracle but defer the details to the next section.

The resulting model is a strict strengthening of the *oblivious algorithms* framework used by Daskalakis-Papadimitriou [DP09] to show optimality of the QPTAS for approximate equilibria (longer discussion in Section 2.3). Our lower bounds thus immediately strengthen theirs.

Finally, we must note that our algorithmic model is strictly weaker than integrality gaps for situations where they make sense. This is because by construction, they can never capture the class of all rounding algorithms for a given convex relaxation. More practically, there are important rounding schemes that are not captured in our framework - for e.g., iterative rounding techniques such as those used for the facility location problem [Li13]. At present, we do not know a clean extension of our framework that can capture such rounding methods.

Known Algorithms.   It is not hard to show that $2^{O(n)}$ queries to the verification oracle (and "0" convex relaxation cost) can find exact equilibria. The QPTAS of [LMM03] implies an algorithm with $n^{O(\log(n))}$ queries to find constant approximate equilibria (and "0" convex relaxation cost). Finally, Harrow-Natarajan-Wu [HNW16] show that the same result can be obtained via an algorithm with $n^{O(\log(n))}$-convex relaxation cost and "0" queries to the verification oracle.

We are now ready to summarize our main hardness results for computing equilibria in this framework.

Summary of Results.   Our main results establish that the known algorithms for computing equilibria, both exact and approximate, are optimal, up to polynomial factors in our model. Specifically, we show that for two-player games with n strategies and all payoffs in $[-1, 1]$,

1. for any $\varepsilon \leqslant \Theta(1/n^4)$, there's no algorithm that uses a $2^{o(n)}$-time SoS relaxation to construct a $2^{o(n)}$-size list of candidates to compute an ε-approximate NE.

2. for some constant $\varepsilon > 0$, there's no algorithm that uses a $n^{o(\log(n))}$-time SoS relaxation to construct a list of size $n^{o(\log(n))}$ to compute an ε-approximate equilibrium.

Our results can be seen as unconditional confirmation, in our restricted algorithmic framework, of the recent result of Rubinstein [Rub16].

Since ours is a stronger model of computation, our lower-bounds strictly improve upon those of Daskalakis-Papadimitriou [DP09]. Along the way, we also strengthen the lower bound in their more

restricted model (see Section 2.2). We also obtain a "gapped" (approximation hardness) version improving on the classical result of Gilboa-Zemel [GZ89] (see Section 2.2).

We give a brief outline of what follows next before the technical sections. In Section 2, we present a formal description of our algorithmic model and state precise versions of our results. We also discuss some of the improvements to previous work that we obtain along the way. In Section 3, we give an overview of the main technical ideas required in our proofs. Finally, in Section 2.3, we briefly discuss important related work.

2 Algorithmic Model and Statements of Results

We first set up the standard terminology for talking about games and equilibria. A game G between Alice and Bob with n strategies for both is described by two payoff matrices R and C in $[-1, 1]^{n \times n}$,[1] where Alice plays rows and Bob plays columns. A mixed-strategy (now on called strategy) for either of the players is an element of $\Delta_n = \{x \in [0, 1]^n \mid \sum_{i=1}^n x_i = 1\}$ and encodes a probability distribution on the space of n pure strategies $\{1, \dots, n\}$ denoted by $[n]$ now on. We write $e_i \in \Delta_n$, for the i^{th} pure strategy. When Alice plays $x \in \Delta_n$ and Bob $y \in \Delta_n$, the expected *payoffs* for the two players are $x^\top R y$ and $x^\top C y$, respectively. A Nash Equilibrium (NE) is a strategy pair (x, y) such that no player gains by deviating unilaterally. That is,

$$x^\top R y \geqslant e_i^\top R y \text{ for all } i \in [n]; \quad x^\top C y \geqslant x^\top C e_j \text{ for all } j \in [n]. \tag{2.1}$$

At an ε-approximate NE (ε-NE) no player gains by more than ε by deviating unilaterally:

$$x^\top R y \geqslant e_i^\top R y - \varepsilon \text{ for all } i \in [n]; \quad x^\top C y \geqslant x^\top C e_j - \varepsilon \text{ for all } j \in [n]. \tag{2.2}$$

Sum-of-Squares Method and *Pseudo-equilibria*. The sum-of-squares method is a sequence of increasingly tight SDP relaxations, indexed by an integer parameter d, for the problem of finding a solution to a system of polynomial inequalities in real valued variables. We provide a brief overview of this method specialized to computing (exact/approximate) equilibria here and point the reader to the lecture notes [BS] for further details on method and its applications.

Central to the sum of squares method is the notion of a pseudo-distribution that generalizes probability distributions.

Definition 2.1 (Pseudo-distribution). A degree d pseudo-distribution is a finitely supported signed measure $\tilde{\mu}$ on $\mathbb{R}^n$ such that the associated linear functional (pseudo-expectation) $\tilde{\mathbb{E}}$ that maps any function $f : \mathbb{R}^n \to \mathbb{R}$ to $\tilde{\mathbb{E}}[f] = \sum_{x : \tilde{\mu}(x) \neq 0} \tilde{\mu}(x) f(x)$ satisfies the following properties:

1. **Normalization:** $\tilde{\mathbb{E}}[1] = 1$ or equivalently, $\sum_{x : \tilde{\mu}(x) \neq 0} \tilde{\mu}(x) = 1$, and

2. **Positivity:** $\tilde{\mathbb{E}}[q^2] \geqslant 0$ for every degree $\leqslant d/2$ polynomial q on $\mathbb{R}^n$.

A pseudo-distribution $\tilde{\mu}$ is said to satisfy a polynomial inequality constraint $q \geqslant 0$ if for every polynomial p, $\tilde{\mathbb{E}}[p^2 q] \geqslant 0$ whenever $\deg(p^2 q) \leqslant d$.

It's not hard to show that any degree ∞ pseudo-distribution is a probability distribution on $\mathbb{R}^n$.

---

[1] Adding dummy strategies to make the number of strategies equal, and normalizing payoffs to lie in $[-1, 1]$ is without loss of generality.

Since $\tilde{\mathbb{E}}_{\tilde{\mu}}$ is a linear operator, it is completely specified by its action on any basis of $\mathcal{P}_d$. The monomial $x^S = \Pi_i x_i^{S_i}$ for any $S \in \mathbb{Z}^n$, $\sum_{i=1}^n S_i \leqslant d$ (the total degree) form a basis for $\mathcal{P}_d$. Let $(1, x) \in \mathbb{R}^{n+1}$ denote the vector with first coordinate 1 and last n coordinates matching $x \in \mathbb{R}^n$. Then, the order-d tensor $\tilde{\mathbb{E}}_{\tilde{\mu}}[(1, x)^{\otimes d}]$ has entries equal to pseudo-moments of $\tilde{\mu}$ and thus completely describes the pseudo-expectation $\tilde{\mathbb{E}}_{\tilde{\mu}}$.

When $\tilde{\mu}$ is a probability measure on $\mathbb{R}^n$, the associated pseudo-expectation is an actual expectation operator. Key to the utility of the notion of pseudo-distributions is the following classical fact.

Fact 2.2 (Lasserre, Shor, Parillo, Nesterov [Las01, NWY00, Par00, Sho89]). *Let $g_1, g_2, \ldots, g_r, h_1, h_2, \ldots, h_q$ be polynomials of degree at most d. Then, the convex set*

$$\mathbb{C}_d = \{\tilde{\mathbb{E}}_{\tilde{\mu}}[(1, x)^{\otimes d}] \mid \tilde{\mu}: \deg d, \text{ satisfies } g_1 = 0, g_2 = 0, \ldots, g_r = 0, h_1 \geqslant 0, h_2 \geqslant 0, \ldots, h_s \geqslant 0\}$$

has an $n^{O(d)}$-time weak separation oracle in the sense of [GLS81].

Definition 2.3 (Sum-of-Squares Algorithm). The degree d sum-of-squares algorithm takes input a system of polynomial constraints $g_1 = 0, g_2 = 0, \ldots, g_r = 0, h_1 \geqslant 0, h_2 \geqslant 0, \ldots, h_s \geqslant 0$ each of degree at most d and either outputs "infeasible" or returns a degree d pseudo-expectation $\tilde{\mathbb{E}}_{\tilde{\mu}}$ satisfying Definition 2.4. Fact 2.2 implies that the degree d SoS algorithm runs in time $n^{O(d)}$.

Pseudo-distributions naturally suggest a relaxation of the notion of equilibrium, which we call, *pseudo-equilibrium*. The advantage is that unlike equilibria, pseudo-equilibria of degree d are *efficiently* computable using Fact 2.2. It might be helpful to note that degree 2 pseudo-equilibria (without the semi-definite constraints) correspond to the well-studied notion of *correlated* equilibria.

Definition 2.4 (Degree d Pseudo-equilibrium). Given a two player game (R, C), a degree d pseudo-equilibrium for (R, C) is a degree d pseudo-distribution on strategy profiles (x, y) satisfying the quadratic polynomial inequality constraints in (2.1). A degree d, ε-approximate pseudo-equilibrium is a degree d pseudo-distribution over strategy profiles (x, y) satisfying the quadratic polynomial constraints in (2.2).

2.1 Rounding Framework for Equilibria

In this section, we present our algorithmic framework based on restricted roundings of sum-of-squares relaxations. While the framework naturally generalizes to any polynomial feasibility problems, we will focus only on (exact) NE here.

Let $\Psi(G)$ be the system of polynomial inequalities in $x, y \in \mathbb{R}^n$ parameterized by a game G represented by $R, C \in \mathbb{R}^{n \times n}$ described in (2.1).

A rounding algorithm takes a solution to the SoS relaxation for $\Psi(G)$ and outputs a candidate equilibrium strategy profile (x^*, y^*). Before presenting our restricted rounding framework, we formally define what we mean by a rounding algorithm.

Definition 2.5 (Rounding Algorithm). A degree d rounding algorithm for $\Psi(G)$ takes input a game G and a degree d pseudo-expectation satisfying the NE constraints of G and outputs a solution to $\Psi(G)$.

Remark 2.6. As discussed, this definition captures all algorithms as one could simply ignore the degree d pseudo-expectation. Notice that we place no restriction on the running time of the rounding algorithm at all.

Next, we present a meaningful restriction of arbitrary roundings that forces the algorithm to "use" the solution to the relaxation. We do this by simply disallowing the access to the game G itself.

Definition 2.7 (Oblivious Rounding). A degree d oblivious rounding algorithm is a rounding algorithm that does not get access to the instance (i.e., game) G itself. The cost of an oblivious rounding algorithm is the degree of the SoS solution it accesses.

Remark 2.8. Two comments are in order:

1. Observe that we do not restrict the running time of an oblivious algorithm. In that respect, this model and the resulting lower bounds are information theoretic.

2. This definition was first proposed in [FFT16] for optimization problems (ours is a "feasibility" problem as it has no objective) who gave a characterization of problems where oblivious rounding algorithms achieved the underlying integrality gap.

While this may appear like a restricted framework, it nonetheless captures several famous algorithms obtained via convex relaxation + rounding paradigm such as those for Max-Cut, Sparsest-Cut, Unique Games etc.

To allow our rounding algorithm more power and capture a longer list of existing algorithms, we allow the oblivious rounding algorithm generate a list of candidate solutions instead of a single one and then check if one of them is indeed a solution using access to a *verification oracle*. This allows us to (i) capture most known rounding algorithms mentioned above with the exception of techniques such as iterative rounding (ii) provide our model the ability to simulate "enumeration" over a restricted search space - a widely used technique in computing equilibria for various classes of games.

To formalize this model, we will allow the algorithm to access the game through a *verification oracle* in addition to a degree d pseudo-equilibrium.

Definition 2.9 (Verification Oracle). A verification oracle for an instance (i.e., game) G takes input a candidate solution (i.e., a strategy profile (x, y)) and correctly outputs "accept" if the candidate is a true solution (i.e, an equilibrium) and "reject" otherwise.

We can now augment the oblivious rounding algorithms with a verification oracle access to the underlying instance (game) G.

Definition 2.10 (Oblivious Rounding with Verification Oracle OV). A degree d, q-query *oblivious rounding* algorithm *with verification oracle* is a degree d oblivious rounding algorithm that, in addition, accessed a verification oracle for the underlying instance G at most q times.

Remark 2.11. Three comments are in order, again.

1. As in the case of oblivious rounding algorithms, note that we do not restrict/measure the running time of the algorithm.

2. Observe that the model allows "adaptivity" - the candidate solutions can be generated after looking at the replies of the verification oracle for previous candidates.

3. This model strictly strengthens the "Oblivious algorithms" model studied in [DP09]. Oblivious algorithms, in the sense of [DP09] are simply degree 0 OV algorithms. That is, OV algorithms that get access only to the verification oracle and neither the game nor the pseudo-equilibrium.

Finally, for the special case of equilibria, we can allow an even stronger access to the underlying game through a more powerful oracle. Observe that if we fix one of x or y in (2.1) and (2.2), then the resulting system of constraints is linear in the other variables. Thus, one can check feasibility of this resulting system in polynomial time.

Motivated by this, we can define an oracle that takes input a "partial strategy" - either one of x or y - and check if there's an equilibrium that agrees with the given partial strategy. This definition, unlike all others, is not meaningful for arbitrary polynomial systems $\Psi(G)$.

Definition 2.12 (Partial Verification Oracle). A partial verification oracle for a game G takes input one of x, y and outputs "accept" if there's an equilibrium for G that agrees with the input and "reject" otherwise.

Remark 2.13. It is important to observe that a partial verification oracle is a *stronger* access to the underlying game - namely, it allows checking if there's an equilibrium matching only one of the player's strategies. Indeed, as we will note in the discussion of our results, the lower bound of Daskalakis-Papadimitriou does *not* hold for algorithms that have this stronger access to game via a Partial verification oracle.

As a by product of our proof, we obtain a strengthening of their result that holds against oblivious algorithms that get access to the stronger partial verification oracle.

2.2 Our Results

Our main results are optimal lower bounds for OV roundings of Sum-of-Squares for computing exact and approximate equilibria. Throughout this section, we refer to a two-player game with n strategies each and all payoffs bounded in $[-1, 1]$ by just a "game".

Our result for computing exact equilibria is actually robust and allows a certain small but non-trivial relaxation of the notion of equilibria. More precisely, we show that:

Theorem 2.14 (Exponential Lower Bound for NE). *Suppose there is degree d, q query OV rounding for finding $\varepsilon = \Theta(1/n^4)$-approximate equilibria in games. Then, either $d = \Omega(n)$ or $q = 2^{\Omega(n)}$. Further, the same results holds for OV algorithms with the stronger partial verification oracle.*

This follows from a construction of a family of games with properties captured in the following theorem:

Theorem 2.15 (Exponentially large hard family). *For every n large enough, there's a family of $\Gamma = 2^{\Theta(n)}$ games $\{G_i = (A_i, B_i)\}_{i=1}^{\Gamma}$ with n pure strategies for both players and all payoffs in $[-1, 1]$ such that:*
*(1.) **Completeness:** There exists a degree $\Theta(n)$, shared pseudo-equilibrium for every G_i simultaneously.*
*(2.) **Soundness:** For any $i \neq j$ if (x, y) and (x', y') are $\Omega(1/n^4)$-NE of games G_i and G_j respectively then $x \neq x'$ and $y \neq y'$.*

To see why this implies Theorem 2.14, choose an input game G uniformly at random from the family described by the Theorem 2.15. Then, any (potentially randomized) algorithm that succeeds with 2/3 probabilty in computing an ε-NE in G requires either $d = \Omega(n)$ or $q = 2^{\Omega(n)}$. This is because, using property (2) above, such an algorithm must uniquely and correctly identify the input game with probability at least 2/3. The SoS relaxation can just output the shared pseudo-equilibrium giving zero information about the input game. Further, each "reject" answer from the (partial) verification oracle rules out exactly one possibility. Thus, the OV algorithm must use $2^{\Omega(n)}$ queries.

In establishing this result, we prove a "gapped" version of the classical hardness reduction of Gilboa-Zemel from k-clique to finding *well-supported* equilibria with payoff at leaset $\delta > 0$. While in a Nash equilibrium, each player's strategy must put a non-zero probability only on the pure strategies with maximum payoffs, in an ε-well-supported NE, the support of any player's strategy can contain pure strategies that have an ε additively smaller payoff compared to the maximum (see Section 3 for a formal definition). Well-supported equilibria are a weaker solution concept than equilibria and thus, this is technically stronger than one that shows a gapped hardness for finding NE.

Theorem 2.16. *There's a polynomial time reduction that takes input a graph H on n vertices and parameters $k \leqslant n$, $\varepsilon = \Omega(1/n^2)$ and produces a two-player game G with $O(n)$ strategies for each player with all payoffs in $[-1, 1]$ such that (i) if H has a k-clique there exists a NE for G with payoff $\delta > 0$, and (ii) if H has no k-Clique then all ε-well-supported NE of G have payoffs at most $(\delta - \varepsilon)$.*

Our second main result is a quasi-polynomial lower bound on OV algorithms for finding constant-approximate equilibria.

Theorem 2.17. *Suppose there exists a degree d, q-query algorithm for computing ε-approximate equilibria in games for some small enough $\varepsilon = \Theta(1)$. Then, either $d = \Omega(\log(n))$ or $q = n^{\Omega(\log(n))}$. Further, the same result continues to hold for OV algorithms that use the stronger partial verification oracle.*

As above, we obtain this theorem by giving a construction of a hard family of games:

Theorem 2.18 (Quasi-polynomial Hardness for $O(1)$-NE). *For every n large enough, there's a family of $\Gamma = n^{\Omega(\log(n))}$ two-player games $\{G_i = (R_i, C_i)\}_{i=1}^{\Gamma}$ with n strategies and all payoffs in $[-1, 1]$ s.t.:*
*(1.) **Completeness:** All G_i's share a common degree $\Theta(\log(n))$ pseudo-equilibrium.*
*(2.) **Soundness:** There exists an $\varepsilon = O(1)$, such that for any i, j and any pair of ε-NE in G_i and G_j, say (x, y) and (x', y') respectively, $x \neq x'$ and $y \neq y'$.*

Along the way to this result, we extend the result of Daskalakis and Papadimitriou on lower bounds for their model of oblivious algorithms. Their construction, in the language of this paper, shows lower bounds on the number of calls to the Verification Oracle 2.9 described before [DP09]. We extend it to lower bound against the stronger Oracle 2.12.

Corollary 2.19. *For a given game (R, C) with n strategies, if an algorithm queries Oracle 2.12 q times and outputs an $O(1)$-NE of (R, C), then $q = n^{\Omega(\log(n))}$.*

Remark 2.20 (Well-Supported NE vs NE). Three points in order.

1. The notion of approximate well-supported NE has been of independent interest and studied as such (see for e.g. [CDT09, DP09]). Since every ε-WNE of a game is also its ε-NE, all our lower bounds hold for finding WNE as well.

2. Our proofs, in fact, yield a stronger hardness result for finding ε-WNE : we show that finding $\varepsilon = O(1/n^2)$-WNE requires an exponential degree or queries to the (partial) verification oracle.

3. One can strengthen our results to an oracle that gets another strong access to the underlying game: given the support of both player's strategies, accept if there's an equilibrium with those supports. Our lower bounds work against such verification oracles too.

2.3   Related Work

We discuss three works that are especially relevant to this work.

The sum-of-squares SDP hierarchy was recently studied in the context of computing equilibria by Harrow-Natarajan-Wu [HNW16]. In the language of the present paper, they showed that (*i*) there's a quasi-polynomial degree oblivious rounding (that does not need access to the verification oracle) that matches the guarantees of the algorithm in [LMM03] and (*ii*) there's a cost[2] measure on equilibria such that the degree $\Omega(\log(n))$ sum-of-squares algorithm has a constant-factor integrality gap in approximating it over the space of all constant approximate equilibria. While both the works show a hardness for the sum-of-squares algorithm for computing equilibria, in contrast to their work that focuses on showing an integrality gap for a certain objective measure over equilibria, we focus on the task of computing *any* equilibrium.

The notion of oblivious roundings was first studied in a work due to Feige, Feldman and Talgam-Cohen [FFT16]. Their work gave a neat characterization of the problems where oblivious roundings can achieve the integrality gap (for optimization, instead of feasibility problems that we study in this work) of the underlying relaxation. Their main application was for the maximum welfare problem where they showed that if the valuation functions of the agents are submodular, then, oblivious rounding achieves the integrality gap of the natural configuration LP relaxation for the problem, while, for the case of gross substitute valuations, oblivious roundings cannot achieve the integrality gap of the configuration LP for the problem.

The work of Daskalakis-Papadimitriou [DP09] studied algorithms that *obliviously* (= degree 0 OV algorithms using the standard verification oracle 2.9) search the smallest possible space of strategy profiles to find an approximate equilibrium. Their main motivation was showing that among all oblivious algorithms, the Lipton-Markakis-Mehta algorithm is almost optimal for computing approximate equilibria in games. Their lower bounds, however do not apply to algorithms that use the stronger partial verification oracle 2.12.

While this can appear somewhat technical, this difference comes from a fairly intuitive difference between our two constructions. The construction of [DP09] gives a collection of $\approx n^{\log(n)}$ games such that every approximate equilibrium of any single game is different from every approximate equilibrium of any other. However, each game admits an approximate equilibria that agree on the strategy of one of the players. While this is enough to fool the standard verification oracle (=oblivious algorithms), it does not suffice to fool the stronger partial verification oracle.

Acknowledgement

---

[2]While the proof in https://dspace.mit.edu/openaccess-disseminate/1721.1/109803 shows a result for a somewhat artificial objective function (see Theorem 3.11 and the proof that follows for the description of this objective), in private communication with the authors, we learned that their techniques extend to taking the objective function as social welfare with some more work

3 Technical Overview

Our main results establish lower bounds that match the existing upper bounds for OV algorithms with the stronge partial verification oracle for computing $\Theta(1/n^4)$ (robust, exact) and $\Theta(1)$-approximate Nash equilibrium computation.

Both our constructions will follow via three modular steps.

1. Construct a family of games that are hard for degree 0 OV algorithms - i.e., algorithms that rely only on the (partial) verification oracle. This will strengthen the result in [DP09] for constant-approximate equilibria and provide a new construction for the case of (exact) $1/poly(n)$-approximate equilibria.

2. Construct a game such that approximating the equilibrium that maximizes the social welfare (sum of the payoffs) of the two players (social welfare) is hard for degree d SoS algorithm. It is important to stress this point: even though our final aim is to construct a lower bound for the feasibility problem, the SoS hardness is for the *optimization* version, namely, for maximizing the social welfare. As discussed in the related work, the proof in [HNW16] does not give such a result at the moment. However, there are hardness results in the optimization setting (we rely on the recent work in [DFS16]) based on standard conjectures such as $P \neq NP$ or Exponential Time Hypothesis in this setting which we leverage in our construction.

3. A black-box result that allows "stitching" together the two kinds of hard games into a single family of games that are hard for OV algorithms. The technical heart of the proof here is establishing that the process of "stitching" together the two kinds of hard games does not introduce new spurious (approximate) equilibria that destroy the first or the second property above.

This overview is organized as follows. First, we describe the ideas in our "stitching" theorem as it applies to both our lower-bound results. This part holds in a fairly general setting. Second, we describe the ideas in the construction of the two kinds of games above to complete the proofs of our two lower bound results.

3.1 Stitching Hard Games Together

Before proceeding, we need to define the notion of well-supported equilibria: a well-studied relaxation of the notion of equilibria [CDT09, DP09] that we briefly discussed in Section 2.2.

For integer n and any $T \subseteq [n]$, let $\mathcal{U}_{n,T}$ be probability vector of a distribution on $[n]$ that is uniform over T.

Definition 3.1 (Well-Supported Equilibria). A strategy profile (x, y) of game with payoff matrices (R, C) is called an ε-approximate well-supported NE (in short, ε-WNE) iff every pure strategy in the support of x has a payoff within ε of the maximum given the column player plays y and vice-versa. That is,

$$\text{for every } i \in [n], \ x_i > 0 \Rightarrow e_i^\top R y \geq \max_{k \in [n]} e_k^\top R y - \varepsilon \quad \text{and} \quad y_i > 0 \Rightarrow x^\top C e_i \geq \max_{k \in [n]} x^\top C e_k - \varepsilon. \ (3.1)$$

Remark 3.2. Observe that any NE is an ε-WNE . And every ε-WNE is an ε-approximate NE. However, an ε-approximate NE, in general, is not an ε-WNE .

Next, we formalize the two kinds of hard games that we discussed in the previous subsection. Our requirement from the SoSHard game is straightforward. We want a game where the social welfare of a pseudo-equilibrium is higher than the social welfare of every ε-WNE by at least 2ε.

Definition 3.3 (degree t, (ε, δ)-SoSHard game). For some $\varepsilon < \delta$, a game (R, C) is said to be degree t, (ε, δ)-SoSHard game if there's a degree t pseudo-equilibrium with payoffs for each player exceeding $\delta > 0$ while the social welfare (sum of the payoffs of the players) in every ε-WNE is at most $2(\delta - \varepsilon)$.

Our requirement from the hard family of games for oblivious algorithms, i.e., enumerative techniques, denoted by EnumHard games, is the following. We want a family of games with K strategies each, that is indexed by a low-intersection family of subsets of $[K]$ such that all ε-WNE of a game indexed by S are close to the uniform distribution on the strategies in $S \subseteq [K]$.

Definition 3.4 ((ε, τ)-EnumHard game). For some $\beta > 0$, $\mathcal{F}$ be a family of subsets of $[K]$ so that for any pair $S, T \in \mathcal{F}$, $|S \Delta T| \geqslant \beta\tau\varepsilon K$, or, equivalently, $\|\mathcal{U}_{K,S} - \mathcal{U}_{K,T}\|_1 \geqslant \beta\tau\varepsilon$. A family of games (R_S, C_S) with K strategies each, indexed by subsets $S \in \mathcal{F}$ is said to be (ε, τ)-EnumHard if any ε-WNE (x, y) of (R_S, C_S) satisfies: $\|x - \mathcal{U}_{K,S}\|_1, \|y - \mathcal{U}_{K,S}\|_1 \leqslant \tau$.

The SoSHard game above gives degree t SoS hardness of obtaining high payoff ε-WNE (and thus, also for finding NE). Similarly, for $\beta > 2$, the family of EnumHard games constructed for each $S \in \mathcal{F}$ would give lower bound of $q(\beta)$-queries to the verification oracle to find ε-WNE . Observe, however, that neither hardness hold for the notion of eqilibria we care about: ε-approximate NE.

Our goal now is to combine the above to constructions to obtain a family of games for computing ε-approximate NE via OV algorithms with appropriate degree and number of queries. The major issue in combining two games is to handle equilibria that are supported on startegies of both the games. The goal of our construction is to avoid these.

Construction of the Hard Family. Let (R, C) be $N \times N$ game and (R_S, C_S) be $K \times K$ game with all payoffs in $[-1, 1]$. We will use these two constructions to obtain a family of games (R'_S, C'_S) with $(N + K)$ pure strategies for each player and show :

1. **SoS Completeness:** any pseudo-equilibrium for the SoSHard game with payoff at least δ for each player is a pseudo-equilibrium for every (R'_S, C'_S), and

2. **Soundness:** every ε-NE of (R'_S, C'_S) is close, in L1 distance, to $\mathcal{U}_S$.

We first describe the construction. Throughout this overview, it will be easier for the sake of exposition to allow payoffs in our games to exceed 1 in magnitude - we finally just have to normalize the game to bring all payoffs in $[-1, 1]$.

For every $S \in \mathcal{F}$, the family of sets $\mathcal{F}$ indexing the EnumHard family of games, we construct a new game (R'_S, C'_S) as follows (where δ and -2 denote blocks of appropriate sizes with all entries equal to δ and -2 respectively).

$$R'_S = \begin{bmatrix} R & -2 \\ \delta & R_S \end{bmatrix} \quad \text{and} \quad C'_S = \begin{bmatrix} C & \delta \\ -2 & C_S \end{bmatrix} \tag{3.2}$$

Completeness. The proof of *completeness* (Lemma 4.4) is intuitively based on the following elementary observation. Suppose (R, S) had an equilibrium (x, y) with payoff at least δ for each player. Then, we note that (x, y) is also an equilibrium in the new game (R'_S, C'_S). This is simply because the row (respectively, column) player has no incentive to use strategies in the lower (respectively, right) block. Of course, (R, S) does not have such an equilibrium, only a pseudo-equilibrium, but it turns out that the same argument as above extends to pseudo-equilibria.

Soundness. The proof of *soundness* requires a lengthier argument. The proof is based on two intermediate claims: (i) every ε-WNE of (R'_S, C'_S) is supported only on last K strategies for both the players (Lemma 4.5) and (ii) every ε-NE of (R'_S, C'_S) is close (in L1 distance) to an ε-WNE . It is not hard to show soundness using these two claims, that is, that each ε-NE of the game is close to (in L1 distance) $(\mathbf{0}_N, \mathcal{U}_{K,S})$ (Lemma 4.7). From the first claim, we can conclude that every ε-WNE of game (R'_S, C'_S) is essentially an ε-WNE of game (R_S, C_S). By the second claim and triangle inequality for L1 distance, $\tau\varepsilon$ close to $\mathcal{U}_{K,S}$. The proof of the second claim above follows from essentially a straightforward calculation and incurs a factor τ loss in the distance and a quadratic loss in the approximation.

To prove the first claim, we first note the following "dichotomy" in the game: either both players play some of first N pure strategies or neither does. This is immediate because of the presence of off-diagonal block with all payoffs (-2) in the two payoff matrices. Next, we observe that if the players' strategy is supported on some of the first N pure strategies, then, in fact, the probability mass on the first N strategies should be $\Theta(1)$ for each. This is again a direct consequence of the presence of the (-2) off-diagnal block in the payoff matrices.

Next, we observe that if we take any ε-WNE of (R'_S, C'_S) and condition on using only the first N strategies, then the resulting strategy pair should be an ε-WNE for (R, C). This is because for any $i \in [N]$, the payoff from playing i given that the second player plays any strategy in the second block is equal and (-2). By the soundness property of SoSHard game, the total payoff from this first $N \times N$ block of the game, thus, can be at most $2(\delta - O(\varepsilon))$. On the other hand, for any strategy in the first $N \times N$ block, each of the last K strategies for the other player give a payoff of at least δ for both players. Thus, there's an incentive for at least one of the players to deviate to playing one of the K strategies. Using the "dichotomy" claim above, the other player must also deviate similarly.

These arguments ultimately yield the following "stiching-together" result:

Theorem 3.5. *Given parameters $\varepsilon, \delta > \varepsilon$, and $\tau > 0$, (i) let (R, C) be a degree t, (ε, δ)-SoSHard game of dimension $N \times N$, and (ii) for an integer K, appropriately chosen β, and subsets $S_1, \ldots, S_{q(\beta)}$ of $[K]$ such that $\|\mathcal{U}_{K,S_i} - \mathcal{U}_{K,S_j}\|_1 > \beta\tau\varepsilon$, let (R_{S_i}, C_{S_i}) be an (ε, τ)-EnumHard game for each $i \in \{1, \ldots, q(\beta)\}$.*

Then, for the family of games $\mathcal{F} = \{\mathcal{G}_i \mid \mathcal{G}_i$ is the game of (3.2) constructed using (R, C) and $(R_{S_i}, C_{S_i})\}$,
(1.) All $\mathcal{G}_i$'s have a common degree t pseudo-equilibrium.
(2.) For any pair of games $\mathcal{G}_i \neq \mathcal{G}_j \in \mathcal{F}$, their $O(\varepsilon^2)$-NE strategy sets of either players do not intersect.

The above theorem implies that if there exists degree d, q-query algorithm to find ε^2-NE in two player games, then $d = \Omega(t)$ and $q = \Omega(q(\beta))$.

3.2 $O(n^{\log(n)})$ lower bound for the constant approximation

Given the discussion in Section 3.1, we only need to give separate constructions of SoSHard games with degree $d = \Theta(\log(n))$, and $n^{\Omega(\log(n))}$-large family of EnumHard games for some constant $\varepsilon > 0$.

Reduction from planted clique does not work, yet. The first candidate here, of course, is the reduction of Hazan-Krauthgamer [HK11] who show such a reduction from the planted clique problem in average case complexity. A recent work of Barak-Hopkins-Kelner-Kothari-Moitra-Potechin [BHK+16] shows a quasi-polynomial SoS lower bound for the planted clique problem. Thus, a priori, it appears that we immediately obtain the SoSHard game as required here.

There's a subtle issue, however, that prevents this program from going through. For a reduction to carry over and give a SoS lower bound, we must show that the map that takes a solution to the starting problem (i.e., planted clique) into a solution to the problem of interest (approximate equilibria maximizing social welfare) be computed by a low-degree polynomial (see Fact 5.6).

Now, in the reduction due to Hazan-Krauthgamer, this map is really simple: start from the planted clique, say S, and return uniform distribution on pure strategies corresponding to the vertices in the clique, say $\mathcal{U}_S$. If x is the indicator vector of the clique, then, the resulting equilibrium is essentially described by $x/|S|$. This appears like a degree 1 map except for one important issue - each non-zero entry of $\mathcal{U}_S$ is $1/|S|$. If the size of the clique we planted is "fixed", then, this entry is a constant and the map is indeed a degree 1 polynomial. However, technically, fixing the size of the clique requires that the SoS lower bound hold for the feasibility version of the polynomial program for clique that has an explicit size-constraint for the clique.

Unfortunately, at present, the lower bound in [BHK+16] holds only for the optimization version of the clique program that doesn't have an explicit clique size constraint. While we fully expect the SoS lower bound to hold for both versions of the clique program, at present there are significant technical difficulties in extending the lower bound to the feasibility version.

Relying on known ETH hardness results. The starting point of our construction of the SoSHard game is the recent work [DFS16] that shows a construction of a game (R, C) via reduction from *free repetitions* of 3SAT (an idea that began in the work of [AIM14] and first used to show hardness of computing equilibria in [BKW14]) to show hardnesss of maximizing social welfare in two-player games modulo the ETH. Now, there are standard hardness results for approximating 3SAT via SoS (i.e., Grigoriev's Theorem). Thus, we will be done if we could show that the reduction appearing in the ETH hardness holds also when restricted to SoS algorithm. Long story short, there's a simple sufficient condition to check in this regard (variant of which appeared first in the work of Tulsiani [Tul09]) and demands only that the map that takes a solution of 3SAT into a solution of the problem at hand be a low-degree polynomial. We give a brief overview of the reduction (see Section 6.1) in [DFS16] and argue why their reduction induces a low-degree polynomial map between solutions. Checking that the appropriate parameters hold gives us our SoSHard game as required.

Such a result might be of independent interest so we state it here in a standalone form.

Lemma 3.6 (SoSHard game for $O(1)$-NE)**.** *There exists a game (R, C) for $R, C \in [-1, 1]^{N \times N}$ such that:*

1. **Completeness:** *There's a degree $\Omega(\log(N))$, pseudo-equilibrium with payoffs $\geqslant 1$ for both players.*

2. **Soundness:** *For every ε-NE (x, y) payoffs of both players is at most $(1 - \varepsilon)$ for $\varepsilon \geqslant 1/1200$.*

The starting point of our construction of EnumHard game is the work of [DP09] discussed before, which we briefly note the important aspects of it before proceeding.

Let $m = \binom{l}{l/2}$ for some even integer l and fix any $\varepsilon > 0$. The family of games (A_S, B_S) is indexed by a collections of subsets $S \subseteq [m]$, each of size ℓ. The key property of this family is that for every ε-WNE of (A_S, B_S), the support of the row player's strategy is contained in S. Further, the row

player's strategy is at most $O(\varepsilon)$ L1 distance away from the $\mathcal{U}_{m,S}$ in the l_1 norm. Finally, for a given $\beta > 0$ there exists a family of at least $n^{(0.8-2\beta\varepsilon)\log(m)}$ many such subsets S, such that for any pair of subsets S, S', $\|\mathcal{U}_{m,S} - \mathcal{U}_{m,S'}\|_1 > \beta\varepsilon$.

As noted before, this construction doesn't give a lower bound against the partial verification oracle as there's an ε-WNE for every game that share the column player's strategy. To construct an EnumHard game, we need such guarantees for both the players.

It appears natural to somehow take an additional flipped copy of the game with the two players' role reversed and combine them somehow to get EnumHard game. The challenge of course is to combine them in such a way that ensures that all constant approximate equilibria are still close to the appropriate uniform strategy. With a slightly technical argument, we show that there's a simple construction that satisfies such a requirement (see Section 6.2).

3.3 Exponential lower bound for the inverse-polynomial approximation

As before, we need to construct a degree $\Omega(n)$ SoSHard game and an exponentially large family of EnumHard games with disjoint set of $O(1/n^2)$-WNE .

For a SoSHard game, as before, we can work with any reduction that (i) starts from a problem that is hard for SoS and (ii) produces a game where approximating the payoff of the best $1/n^2$-approximate WNE is hard.

Our idea for constructing the SoSHard game is again based on a reduction from a problem with known, exponential SoS lower bounds. Specifically, we give a reduction from the independent-set problem to a gap version of approximate Nash equilibrium. Given a graph $G = (V, E)$ on n vertices and a parameter k, we construct a $(2n + 1) \times (2n + 1)$ game (R, C). We show that: (i) if G has a k sized independent set then game (R, C) has an equilibrium with payoff $(1 + 1/2k)$ (Lemma 5.4), (ii) if G does not have an independent set of size k then all $(1/5k)$-WNE of game (R, C) gives payoff at most 1 (Lemma 5.5).

Our construction is inspired by the classical work of Gilboa and Zemel [GZ89] (see (5.1)).

The completeness of this construction is easy. Given an independent set of G, it can be argued that uniform distribution on the support of the independent set is in fact an equilibrium for the game above. Further, the size of the independent set, k, is fixed as a parameter (equivalently, appears as a constraint in the polynomial program we consider). This implies that the equilibrium is computed via a degree 1 (linear) map and we can immediately apply the technology for importing reductions into SoS framework (Fact 5.6). Combining with the known SoS lower bounds for finding independent sets of fixed $k = \Theta(n)$ size [Tul09] completes the proof.

To construct an EnumHard game, we use generalized matching pennies (GMP) game. Here the row-player wants to match strategies with the column-player, while column-player does not want to match, *i.e.*, game (R, C) such that $R(i, i) = 1, C(i, i) = -1$, and the rest are zero. It is easy to see that the uniform distribution over all strategies is the only NE strategy for both the players. We show that in an $m \times m$ game, even $(1/m^2)$-WNE is at distance at most $O(1/m)$ from the uniform strategy in l_1 norm (Lemma 5.8).

Now for every subset $S \subseteq [n]$ we construct an $n \times n$ game (R_S, C_S), and embed $m = |S|$ sized GMP into its $S \times S$ block. We make sure that all strategies outside S are strictly dominated by those inside S by a large margin in payoffs. Thus ε-WNE of game (R_S, C_S) are essentially ε-WNE of the GMP game, and therefore they concentrate around uniform distribution on S, namely $\mathcal{U}_{n,S}$. This gives an EnumHard game (Lemma 5.9). Finally, the property that $\|\mathcal{U}_{n,S} - \mathcal{U}_{n,S'}\|_1 \geqslant \frac{1}{n}$ for any two distinct subsets S, S' of $[n]$, gives us the desired property that the ε-WNE sets of corresponding

14

EnumHard games (R_S, C_S) and $(R_{S'}, C_{S'})$ do not intersect in either player's strategy. This gives us a family of $2^{\Omega(n)}$ many EnumHard games, corresponding to subsets of $[n]$.

The above construction of degree $\Omega(n)$ SoSHard game, and a family of $2^{\Omega(n)}$ many EnumHard games, together with Theorem 3.5 gives Theorem 5.1.

4 Stitching **EnumHard** and **SoSHard** games together

In this section, we show how to combine hardness against SoS and hardness against "enumeration" to obtain lower bounds against OV algorithms. For this we will first work with a stronger notion of approximation, namely ε-well-supported NE (WNE), described next.

4.1 Well-Supported NE vs. Approximate NE

Recall Definition 3.1 of ε-approximate well-supported Nash equilibrium. We state the condition formally below for the reader: For a two player game (R, C), a strategy profile (x, y) is an ε-WNE iff

$$\text{for every } i \leq n, \ x_i > 0 \Rightarrow e_i^\top R y \geq \max_k e_k^\top R y - \varepsilon \quad \text{and} \quad y_i > 0 \Rightarrow x^\top C e_i \geq \max_k x^\top C e_k - \varepsilon.$$

Strategies giving maximum payoff to a player are called her best response, and those giving at least maximum minus ε payoff ε-best response, defined below:

$$\mathsf{BR}_1(y) = \{i \in [n] \mid i \in \text{argmax}_k\, e_k^\top R y\} \quad \text{and} \quad \mathsf{BR}_2(x) = \{i \in [n] \mid i \in \text{argmax}_k\, x^\top C e_i\}$$
$$\varepsilon\text{-}\mathsf{BR}_1(y) = \{i \in [n] \mid e_i^\top R y \geq \max_k e_k^\top R y - \varepsilon\} \quad \text{and} \quad \varepsilon\text{-}\mathsf{BR}_2(x) = \{i \in [n] \mid x^\top C e_i \geq \max_k x^\top C e_k - \varepsilon\}$$

Note that by definition, (x, y) is an ε-WNE iff $\text{supp}(x) \subseteq \varepsilon\text{-}\mathsf{BR}_1(y)$ and $\text{supp}(y) \subseteq \varepsilon\text{-}\mathsf{BR}_2(x)$. The next lemma characterizes NE, approximate NE, WNE , and relates them.

Lemma 4.1. *For any two player game (R, C) with payoffs in $[0, 1]$ the following holds.*

1. *Strategy profile (x, y) is an NE iff $\text{supp}(x) \subseteq \mathsf{BR}_1(y)$ and $\text{supp}(y) \subseteq \mathsf{BR}_2(x)$.*

2. *Every ε-WNE is ε-NE.*

3. *If (x, y) is an ε-NE then $\sum_{i \notin \sqrt{\varepsilon}\text{-}\mathsf{BR}_1(y)} x_i \leq \sqrt{\varepsilon}$ and $\sum_{i \notin \sqrt{\varepsilon}\text{-}\mathsf{BR}_2(x)} y_i \leq \sqrt{\varepsilon}$.*

4. *Given ε-NE (x, y) we can construct $(3\sqrt{\varepsilon})$-WNE (x', y') such that $\text{supp}(x') \subseteq \sqrt{\varepsilon}\text{-}\mathsf{BR}_1(y)$ and $\text{supp}(y') \subseteq \sqrt{\varepsilon}\text{-}\mathsf{BR}_2(x)$.*

Proof. Proofs of first and second part follow by definition. The argument for the third and fourth part described next are essentially from Lemma 2.2 of [CDT09] due to Chen, Deng, and Teng.

For the third part, let $i^* \in \mathsf{BR}_1(y)$. Then, for each $i \notin \sqrt{\varepsilon}\text{-}\mathsf{BR}_1(y)$ we have $e_i^\top R y < e_{i^*}^\top R y - \sqrt{\varepsilon}$. Therefore, if $\sum_{i \notin \sqrt{\varepsilon}\text{-}\mathsf{BR}_1(y)} x_i > \sqrt{\varepsilon}$, then row player's payoff $x^\top R y < e_{i^*}^\top R y - \varepsilon$, a contradiction. Similarly for y.

For the fourth part, construct of (x', y') from (x, y) by removing the probability mass from strategies outside $\sqrt{\varepsilon}\text{-}\mathsf{BR}_1(y)$ and uniformly distribute them on strategies of $\sqrt{\varepsilon}\text{-}\mathsf{BR}_1(y)$. Since payoffs are in $[0, 1]$, this will increase or decrease payoffs of any strategy at most by $\sqrt{\varepsilon}$. Therefore, (x', y') is $(3\sqrt{\varepsilon})$-WNE . And $\text{supp}(x') \subseteq \sqrt{\varepsilon}\text{-}\mathsf{BR}_1(y)$ and $\text{supp}(y') \subseteq \sqrt{\varepsilon}\text{-}\mathsf{BR}_2(x)$. $\qquad\square$

4.2 Abstract Game Construction

Our construction of the hard games are based on combining construction of integrality gap against Sum-of-Squares relaxations computing equilibria that maximize social welfare and games that are hard for enumeration based algorithms. We abstract out the properties of these two kinds of hard games in this section. Our hardness results in the next two sections will invoke the general combining strategy from this section.

First, we define two player games such that maximizing social welfare is hard for SoS.

Definition 4.2 (Games Hard for SoS). A game (R, C) for $R, C \in [-1, 1]^{N \times N}$ is said to be degree t, (ε, δ)-SoSHard for $\varepsilon, \delta \in [0, 1]$ if:

1. **Completeness:** There's a degree t pseudo-equilibrium for the game (R, C) that has payoffs for both players at least δ.

2. **Soundness:** Every ε-WNE (x, y) of (R, C) has total payoff $x^\top (R + C) y < 2(\delta - \varepsilon)$.

Next we need a game to construct the family of games with disjoint set of equilibria, and the property we need from these is as follows.

Definition 4.3 (Games Hard for Enumeration). A $K \times K$ game (R_S, C_S), where $S \subseteq [K]$, is said to be (ε, τ, c)-EnumHard for $\tau > 0$ and $\varepsilon, c \in [0, 1]$ such that cK is an integer, if $R_S, C_S \in [-1, 1]^{K \times K}$, and

1. For any strategy y of the column player $\exists i \in [K]$ such that $e_i^\top R_S y \geqslant 1/2$, and for any x of the row player $\exists j \in [K]$ such that $x C_S e_j \geqslant 1/2$.

2. For any ε-WNE (x, y), let $\tilde{x} = (x_1, \ldots, x_{cK})$ and $\tilde{y} = (y_1, \ldots, y_{cK})$. The following holds:

   (a) $\mathsf{supp}(\tilde{x}), \mathsf{supp}(\tilde{y}) \subseteq S$.
   (b) $\|\tilde{x} - c\mathcal{U}_{cK,S}\|_1 \leqslant \tau\varepsilon$ and $\|\tilde{y} - c\mathcal{U}_{cK,S}\|_1 \leqslant \tau\varepsilon$.

Parameter c in the above definition is needed for the quasi-polynomial lower bound. It is to essentially imply that approximate NE strategy sets, while projected on first cK coordinates, are disjoint and therefore they are disjoint before projection as well.

Given parameters $1/2 \geqslant \delta > \varepsilon > 0$, $c \in [0, 1]$, and $\tau > 0$, let (R, C) be an $N \times N$ game that is degree t, (ε, δ)-SoSHard . And for $K = \mathrm{poly}(n)$ and $S \subset [K]$ let (R_S, C_S) be an (ε, τ, c)-EnumHard game of size $K \times K$. Construct game (R', C') of size $(N + K) \times (N + K)$ as follows:

$$R' = \begin{bmatrix} R & -2 \\ \delta & R_S \end{bmatrix} \quad \text{and} \quad C' = \begin{bmatrix} C & \delta \\ -2 & C_S \end{bmatrix} \tag{4.1}$$

where the constants, namely δ and (-2), represent matrices of appropriate size with all entries set to the stated constant.

Lemma 4.4 (Completeness). *There's degree t pseudo-equilibrium in the game (R', C') that gives a payoff of 1 to both the players.*

Proof. By construction of (R', C') together with property (1) of a degree t, (ε, δ)-SoSHard game, we know that there's a pseudo-distribution of degree t on strategy profiles (x, y) that form an NE for (R, C). This immediately yields a pseudo-distribution $\tilde{\mathbb{E}}$ on strategy profiles for the game (R', C') by padding x and y with zeros on the last K coordinates. We claim that this is a pseudo-distribution on NE of (R', C').

16

Since $\tilde{\mathbb{E}}$ is a pseudo-distribution on NE of (R, C), it satisfies the inequality constraints $x^\top R' y \geqslant e_i^\top R' y$ for every $i \leqslant N$. Further, by property (1), the $\tilde{\mathbb{E}}$ satisfies the constraint $x^\top R' y \geqslant \delta$. On the other hand, by construction of (R', C'), $e_i^\top R' y \leqslant \delta$ since y satisfies $y_i = 0$ for all $i \geqslant (N+1)$. A similar argument shows that $\tilde{\mathbb{E}}$ satisfies the inequality constraint $x^\top C' y \geqslant x^\top C' e_i$ for every $i \leqslant (N + K)$. This yields that $\tilde{\mathbb{E}}$ satisfies the constraints of NE for (R', C') as required. □

To prove the soundness of the construction, we first need to show that every WNE of game (R', C') is supported on the strategies of the EnumHard game.

Lemma 4.5. *For $0 < \varepsilon' < 2\varepsilon/3$, every ε'-WNE (x, y) of game (R', C'), we have $\mathsf{supp}(x), \mathsf{supp}(y) \subset \{(N + 1), \ldots, (N + K)\}$.*

Proof. Let us denote $x_L = \sum_{i \in [N]} x_i$, $x_R = (1 - x_L)$, $y_L = \sum_{i \in [N]} y_i$ and $y_R = (1 - Y_L)$. It suffices to show that $x_L = y_L = 0$. Suppose not, and $x_L > 0$ or $y_L > 0$. First we show that $x_L > 0$ if and only if $y_L > 0$. By construction, if $y_L = 0$ then payoff of the row player from any $i \in [N]$ is (-2) while from any $i > N$ it is $\delta > 0$. In that case if x_i is non-zero for an $i \in [N]$ then it contradicts (x, y) being ε-WNE . Therefore, $x_L = 0$. Similarly, $x_L = 0 \Rightarrow y_L = 0$ follows.

Next we show that if $x_L > 0$ then it should be significantly bigger.

Claim 4.6. If $x_L > 0$ then $x_L > 2/3$. And if $y_L > 0$ then $y_L > 2/3$.

Proof. Let us prove the second part and the first part follows similarly. By property (1) of EnumHard -game (R_S, C_S) there exists a $j \in \{(N + 1), \ldots, (N + K)\}$ such that payoff of the row player from strategy j, namely $e_j^\top R' y$, is at least $\delta y_L + \frac{y_R}{2}$. Therefore, if $x_L > 0$ the strategy $i \in [N]$ with $x_i > 0$ should have payoff at least $\delta y_L + \frac{y_R}{2} - \varepsilon'$. On the other hand, payoff from any $i \leqslant N$ is at most $y_L - 2y_R$, implying.

$$
\begin{aligned}
y_L - 2y_R \geqslant \delta y_L + \tfrac{1}{2} y_R - \varepsilon' \ &\Rightarrow\ y_L(1 - \delta) \geqslant 2.5 y_R - \varepsilon \quad (\text{because } \varepsilon > \varepsilon') \\
&\Rightarrow\ y_L(1 - \delta) \geqslant 2.5 - 2.5 y_L - \delta \quad (\text{because } \delta \geqslant \varepsilon) \\
&\Rightarrow\ y_L(3.5 - \delta) \geqslant 2.5 - \delta \Rightarrow y_L \geqslant 2/3 \quad (\text{because } \delta \leqslant 1/2)
\end{aligned}
$$
□

Let $\hat{x}$ and $\hat{y}$ be truncation of x and y to first N coordinates, and let $\hat{x}' = \hat{x}/x_L$ and $\hat{y}' = \hat{y}/y_L$ be their normalized vectors. Then, row player's payoff from $i \leqslant N$ is $e_i^\top R \hat{y} - 2y_R$. Furthermore, if $x_i > 0$, then this payoff must be at least $\max_{k \leqslant N} e_k^\top R \hat{y} - 2y_R - \varepsilon' \geqslant \max_{k \leqslant N} e_k^\top R \hat{y} - 2y_R - 2\varepsilon/3$. This implies, $\forall i \in [N]$, $x_i > 0 \Rightarrow e_i^\top R \hat{y}/y_L \geqslant \frac{1}{y_L}(\max_{k \leqslant N} e_k^\top R \hat{y} - 2\varepsilon/3) \Rightarrow e_i^\top R \hat{y}' \geqslant \max_{k \leqslant N} e_k^\top R \hat{y}' - \varepsilon$. The last implication uses $y_L > 2/3$ from Claim 4.6. Similarly, we can show that if $\forall i \in [N]$, $y_i > 0 \Rightarrow \hat{x}'^\top C e_i \geqslant \max_{k \in [N} \hat{x}'^\top C e_i - \varepsilon$. These two together implies that $(\hat{x}', \hat{y}')$ is an ε-WNE of game (R, C). Therefore, by property (2) of SoSHard games, $\hat{x}'^\top (R + C) \hat{y}' \leqslant 2(\delta - \varepsilon) \Rightarrow \hat{x}^\top (R + C) \hat{y} \leqslant (\delta - \varepsilon) x_L y_L$.

We have either $\hat{x}^\top R \hat{y} < (\delta - \varepsilon) x_L y_L$ or $\hat{x}^\top C \hat{y} < (1 - \hat{\varepsilon}) x_L y_L$. We will show contradiction for the former and the contradiction for the latter follows similarly by symmetry of the construction of game (R', C').

Suppose, $\hat{x}^\top R \hat{y} < (\delta - \varepsilon) x_L y_L$. There exists $i \in [N]$ such that $x_i > 0$ and $e_i^\top R \hat{y} < (\delta - \varepsilon) y_L$. Therefore, $e_i^\top R' y = e_i^\top R \hat{y} - 2y_R < (\delta - \varepsilon) y_L - 2y_R$. By property (1) of EnumHard game (R_S, C_S), exists $j > N$ such that $e_j^\top R' y \geqslant \delta y_L + \frac{y_R}{2}$. By requirement of ε'-WNE , payoff from strategy i above should be at least payoff from j minus ε'. This gives,

$$
(\delta - \varepsilon) y_L - 2y_R > \delta y_L + \frac{y_R}{2} - \varepsilon' > \delta y_L + \frac{y_R}{2} - \varepsilon \Rightarrow (1 - y_L)\varepsilon > 2.5 y_R \Rightarrow \varepsilon > 2.5
$$

A contradiction. □

Lemma 4.7 (Soundness). *For $0 \leqslant \hat{\varepsilon} < \varepsilon/15$ every $\hat{\varepsilon}^2$-NE (x, y) of game (R', C') satisfies $\|\tilde{x} - (\mathbf{0}_N, c\mathcal{U}_{cK,S})\|_1, \|\tilde{y} - (\mathbf{0}_N, c\mathcal{U}_{cK,S})\|_1 \leqslant 9(\tau + 1)\hat{\varepsilon}$, where $\tilde{x}$ and $\tilde{y}$ are projection of x and y respectively on first $(N + cK)$ coordinates, and $\mathbf{0}_N$ is an N dimensional zero vector.*

Proof. Note that additive scaling of payoff matrices preserves approximate NE, while multiplicative scaling scales the approximation factor by the same amount. Since every entry in R', C' are in $[-2, 1]$ we first add 2 to all of them and then divide them by 3. Thus, (x, y) is now $\hat{\varepsilon}^2/3$-NE of the scaled game, which is also its $\hat{\varepsilon}^2$-NE. Lets denote the scaled game by $(\hat{R}, \hat{C})$.

Then by Lemma 4.1 property (4) there exists (x', y') that is $3\hat{\varepsilon}$-WNE of game $(\hat{R}, \hat{C})$, which is $9\hat{\varepsilon}$-WNE of game (R', C'). Note that $9\hat{\varepsilon} < 2\varepsilon/3$. Applying Lemma 4.5 we get that $\text{supp}(x'), \text{supp}(y') \subseteq \{(N + 1), \ldots, (N + K)\}$. Therefore, strategy (x', y') projected to last K coordinates is a $9\hat{\varepsilon}$-WNE of game (R_S, C_S). The latter is an (ε, τ, c)-EnumHard game, whose property (2.a) implies $\|\tilde{x}' - (\mathbf{0}_N, c\mathcal{U}_{cK,S})\|_1 \leqslant 9\tau\hat{\varepsilon}$, where $\tilde{x}'$ is the projection of $\tilde{x}'$ on first $(N + cK)$ coordinates.

If we get an upper bound on l_1 distance between x and x', then using triangle inequality we will get an upper bound on l_1 distance between $\tilde{x}$ and $c\mathcal{U}_{cK,S}$. By construction of (x', y'), we have that $\text{supp}(x') \subseteq \hat{\varepsilon}\text{-BR}_1(y)$ and $\text{supp}(y') \subseteq \hat{\varepsilon}\text{-BR}_2(x)$. Furthermore, Part (3) of Lemma 4.1 implies $\sum_{i \notin \hat{\varepsilon}\text{-BR}_1(y)} x_i \leqslant \hat{\varepsilon}$ and $\sum_{i \notin \hat{\varepsilon}\text{-BR}_2(x)} y_i \leqslant \hat{\varepsilon}$. While constructing x' from x we set the probability of all strategies not in $\hat{\varepsilon}\text{-BR}_1(y)$ to zero and redistribute their mass to strategies in $\hat{\varepsilon}\text{-BR}_1(y)$ uniformly. Therefore, $\|x - x'\|_1 \leqslant 2\hat{\varepsilon}$. Finally, $\|\tilde{x} - (\mathbf{0}_N, c\mathcal{U}_{cK,S})\|_1 \leqslant \|x - x'\|_1 + \|\tilde{x}' - (\mathbf{0}_N, c\mathcal{U}_{cK,S})\|_1 \leqslant (9\tau + 2)\hat{\varepsilon}$. $\square$

Using the SoSHard game and a family of EnumHard games on far-apart sets, next we construct a family of hard games. The proof of the next theorem will crucially use Lemmas 4.7 and 4.4.

Theorem 4.8. *Given parameters $\varepsilon \in [0, 1/2)$, $\delta > \varepsilon$, $\tau > 0$, and $c \in (0, 1]$*

(h_1) *let (R, C) be a degree t, (ε, δ)-**SoSHard** game of dimension $N \times N$.*

(h_2) *and for $K = \text{poly}(N)$ and subsets $S_1, \ldots, S_f$ of $[K]$ such that $\|\mathcal{U}_{K,S_i} - \mathcal{U}_{K,S_j}\|_1 > \frac{18}{c}(\tau + 1)\varepsilon$, let (R_{S_i}, C_{S_i}) is a (ε, τ, c)-**EnumHard** game for each $i \in \{1, \ldots, f\}$.*

Then, for the game family $\mathcal{F} = \{G_i \mid G_i$ is the game of (4.1) constructed using (R, C) and $(R_{S_i}, C_{S_i})\}$,

1. *All $G_i \in \mathcal{F}$ share a common pseudo-equilibrium of degree t.*

2. *For any pair of games $G_i \neq G_j \in \mathcal{F}$, their set of ε'-NE strategies of either players do not intersect, for all $0 \leqslant \varepsilon' \leqslant (\varepsilon/15)^2$.*

Proof. The first part for each game G_i follows from the first part of Lemma 4.4. For the second part, let (x^i, y^i) and (x^j, y^j) be ε'-NE of games G_i and G_j respectively. If we think of $\varepsilon' = \hat{\varepsilon}^2$ then clearly $\hat{\varepsilon} < \varepsilon/15$ is satisfied. Let $\tilde{x}^i, \tilde{y}^i, \tilde{x}^j, \tilde{y}^j$ be projection of respectively x^i, y^i, x^j, y^j on first $(N + cK)$ coordinates. Using Lemma 4.7 we have $\|\tilde{x}^i - (\mathbf{0}_N, c\mathcal{U}_{cK,S_i})\|_1 \leqslant 9(\tau + 1)$ and $\|\tilde{x}^j - (\mathbf{0}_N, c\mathcal{U}_{cK,S_j})\|_1 \leqslant 9(\tau + 1)$. If $x^i = x^j$ then $\tilde{x}^i = \tilde{x}^j$, and thereby using triangle inequality of l_1 norm we get $\|\mathcal{U}_{cK,S_i} - \mathcal{U}_{cK,S_j}\|_1 \leqslant \frac{18}{c}(\tau + 1)\varepsilon$, a contradiction to hypothesis (h_2). The same contradiction follows if $y^i = y^j$. $\square$

5 Exponential Lower Bounds for NE

In order to prove hardness against using SoS relaxation together with enumeration, we need to construct a family of games all of which share a common high degree pseudo-expectation for its

ε-NE formulation 2.2, while their ε-NE strategy sets for either player are disjoint. In this section, we obtain exponential hardness for $1/\text{poly}(n)$-NE, and prove Theorem 2.14. It follows from the following result.

Theorem 5.1 (Exponential Lower Bound for NE). *For every n large enough, there's a family of $\Gamma = 2^{\Omega(n)}$ two-player games $\{G_i = (R_i, C_i)\}_{i=1}^{\Gamma}$ with $O(n)$ pure strategies for both players and all payoffs in $[-1, 1]$ such that:*

1. **Completeness:** *All G_is share a degree $\Theta(n)$, pseudo-equilibrium.*

2. **Soundness:** *For any $i \neq j$ if (x, y) and (x', y') are $\Theta(1/n^4)$-NE of games G_i and G_j respectively then $x \neq x'$ and $y \neq y'$.*

The proof of Theorem 5.1 follows immediately from the following two constructions of hard games along with the generic combining construction from Theorem 4.8.

Lemma 5.2 (Games Hard for SoS). *There exists a degree $\Theta(n)$, (ε, δ)-SoSHard game (R, C) for $\varepsilon = O(1/n^2)$, $\delta \in (0, 1/2]$ and $R, C \in [-1, 1]^{O(n) \times O(n)}$.*

Lemma 5.3 (Games Hard for Enumeration). *For $K = O(n)$, $\varepsilon = O(1/n^2)$, $\tau = O(n)$ and $c = 1$, there exists a family of subsets $\mathcal{F} \subset 2^{\{0,1\}^K}$ of size at least $2^{\Omega(n)}$ such that*

- *for each $S \in \mathcal{F}$ there is an (ε, τ, c)-EnumHard game (R_S, C_S) for $R_S, C_S \in [-1, 1]^{K \times K}$.*

- *for any pair of distinct $S, S' \in \mathcal{F}$, $\|\mathcal{U}_{cK,S} - \mathcal{U}_{cK,S'}\| > \frac{18}{c}(\tau + 1)\varepsilon$.*

5.1 SoSHard Game Construction

In Lemma 5.2, the SoSHard game has to encode degree $\Theta(n)$ pseudo-equilibrium with high payoff, while no such approximate NE exists. Among the combinatorial problems, such SoS lower bound is known for independent set. If we construct a game that encodes independent set in a NE with high payoffs such that map from the independent set to corresponding NE is low degree then we are done using Fact 5.6. This is exactly what we achieve in this section.

Given an undirected graph G on n vertices represented by an adjacency matrix E (abuse of notation), we wish to check if it has an independent set of size $k \leqslant n$ (k-IS). We will reduce this problem first to finding a Nash equilibrium in a two-player game with certain property. In fact we will show that the game has a NE with payoff at least $\delta = (1 + 1/2k)$ if k-IS has a solution, otherwise all it's ε-WNE have payoff at most $(\delta - \varepsilon)$ for an inverse polynomial ε.

For $\gamma = 1/2$, $A = \mathbf{1}_{n \times n} - E + \gamma I_{n \times n}$, $B = -M * I_{n \times n}$ and $B' = (k + \gamma) * I_{n \times n}$, consider the following $(2n + 1) \times (2n + 1)$ block matrices:

$$R = \begin{bmatrix} A & B & (-1)\mathbf{1}_{2n \times 1} \\ B' & \mathbf{0}_{n \times n} & \\ (1 + \frac{\gamma}{k})\mathbf{1}_{1 \times 2n} & 1 \end{bmatrix} \quad C = \begin{bmatrix} A & B' & (1 + \frac{\gamma}{k})\mathbf{1}_{2n \times 1} \\ B & \mathbf{0}_{n \times n} & \\ (-1)\mathbf{1}_{2n \times 1} & 1 \end{bmatrix} \quad (5.1)$$

In the above construction, submatrix A has entry one in (u, v)th position if edge (u, v) is *not present* in graph G, otherwise the entry is zero. Further, all its diagonals are $(1 + \gamma) = 1.5$. Note that game (R, C) is symmetric, i.e., $C = R^T$. For ease of notation let $m = 2n + 1$.

Lemma 5.4. *If graph G has an independent set $S \subseteq V$ of size k then strategy profile $(\mathbf{x}, \mathbf{x})$ where $\mathbf{x} = (\mathcal{U}_{n,S}, \mathbf{0}_{(n+1)})$ is a Nash equilibrium of game (R, C). Furthermore, payoffs of both the players at this NE is $(1 + 1/2k)$.*

Proof. This is because payoff from each non-zero strategy is $(1 + \gamma/k)$, and from strategies with zero probability it is at most $(1 + \gamma/k)$. □

The above lemma constructs a symmetric NE with high payoff. In what follows, we consider non-existence of both symmetric and non-symmetric NE with high payoff when k-sized independent set does not exist in the graph.

Lemma 5.5. *Given that graph G does not have an independent set of size k, if $(\mathbf{x}, \mathbf{y})$ is an ε-WNE of game (R, C) for $\varepsilon \leqslant \frac{1}{5k}$, then $\mathsf{supp}(\mathbf{x}) = \mathsf{supp}(\mathbf{y}) = \{(2n + 1)\}$.*

Proof. To the contrary suppose $\mathsf{supp}(\mathbf{y}) \neq \{(2n+1)\}$. This proof is based on a number of observations which we list next with a brief justification.

1. $\mathsf{supp}(\mathbf{y}) \cap \{1, \ldots, n\} \neq \emptyset$ and $\mathsf{supp}(\mathbf{x}) \cap \{1, \ldots, n\} \neq \emptyset$: If $\mathsf{supp}(\mathbf{y}) \subseteq \{(n + 1), \ldots, (2n + 1)\}$, then $\mathsf{supp}(\mathbf{x}) = \{(2n + 1)\}$ since any strategy other than the last one gives non-positive payoff to player one, while the last one gives payoff of at least 1. In that case, by the same argument $\mathsf{supp}(\mathbf{y}) = \{(2n + 1)\}$, a contradiction. Second part follows similarly.

2. $\forall i \leqslant n, x_i > 0 \Rightarrow y_i \geqslant \frac{1}{k} - \frac{\varepsilon}{\gamma}$, and $y_i > 0 \Rightarrow x_i \geqslant \frac{1}{k} - \frac{\varepsilon}{\gamma}$: For the first part, suppose $x_i > 0$. Row player's payoff from ith strategy is $e_i^\top R\mathbf{y} \leqslant 1 + \gamma y_i - y_m$. While mth strategy gives her $e_m^\top R\mathbf{y} \geqslant 1 + \gamma/k(1 - y_m)$. Since $(\mathbf{x}, \mathbf{y})$ is ε-WNE , the former payoff has to be at least the latter payoff minus ε,

$$1 + \gamma y_i - y_m \geqslant 1 + \frac{\gamma}{k}(1 - y_m) - \varepsilon \Rightarrow \gamma y_i \geqslant \frac{\gamma}{k} + y_m(1 - \frac{1}{k}) - \varepsilon \Rightarrow y_i \geqslant \frac{1}{k} - \frac{\varepsilon}{\gamma}$$

3. $\mathsf{supp}(\mathbf{x}) \cap \{1, \ldots, n\} = \mathsf{supp}(\mathbf{y}) \cap \{1, \ldots, n\}$ and $|\mathsf{supp}(\mathbf{x}) \cap \{1, \ldots, n\}| \leqslant (k - 1)$: From Point (2), $\mathsf{supp}(\mathbf{x}) \cap \{1, \ldots, n\} = \mathsf{supp}(\mathbf{y}) \cap \{1, \ldots, n\}$ follows. For the second part, to the contrary suppose $|\mathsf{supp}(\mathbf{x}) \cap \{1, \ldots, n\}| \geqslant k$. Since G has no independent set of size k, there exists $u \neq v \in \mathsf{supp}(\mathbf{x}) \cap \{1, \ldots, n\}$ such that edge (u, v) is present in G and in turn $A(u, v) = 0$. Sum of the payoffs of the row player from these two strategies are:

$$e_u^\top R\mathbf{y} + e_v^\top R\mathbf{y} \leqslant (2 - y_u - y_v) + \gamma(y_u + y_v) = 2 - \frac{(y_u + y_v)}{2} \leqslant 2 - \frac{1}{k} + 2\varepsilon$$

The last inequality uses the fact that from Point (2) that if $x_u, x_v > 0$ then $y_u, y_v \geqslant \frac{1}{k} - \frac{\varepsilon}{\gamma}$ and $\gamma = 1/2$. Thus either strategy u or strategy v gives payoff at most $(1 - \frac{1}{2k} + \varepsilon)$. While the payoff from mth strategy is at least 1. Then, ε-WNE ensures that $(1 - \frac{1}{2k} + \varepsilon) \geqslant 1 - \varepsilon \Rightarrow \varepsilon \geqslant \frac{1}{4k}$, a contradiction to $\varepsilon \leqslant \frac{1}{5k}$.

Let $y_L = \sum_{i \in [n]} y_i$ and $a = |\mathsf{supp}(\mathbf{y}) \cap \{1, \ldots, n\}|$ Then from Point (1) we have $y_L > 0$ and from Point (3) we have $a \leqslant (k - 1)$. There exists strategies $i, i' \in \mathsf{supp}(\mathbf{y}) \cap \{1, \ldots, n\}$ such that $y_i \leqslant y_L/a$ and $y_{i'} \geqslant y_L/a$. Again using Point (3) we have $x_i > 0$ as well, and the corresponding payoff is $e_i^\top R\mathbf{y} \leqslant y_L + \gamma y_i - y_m \leqslant y_L + \gamma y_i \leqslant (1 + \gamma/a)y_L$, while the row player's payoff from $(n + i')^{th}$ strategy is $e_{(n+i')}^\top R\mathbf{y} = (k + \gamma)y_{i'} \geqslant (k + \gamma)y_L/a$. Therefore,

$$e_{(n+i')}^\top R\mathbf{y} - e_i^\top R\mathbf{y} \geqslant (k/a + \gamma/a - 1 - \gamma/a)y_L = (k - a)y_L/a \geqslant y_L/a \geqslant \frac{a(\min_{i \in [n]} y_i)}{a} \geqslant \frac{1}{k} - 2\varepsilon > \varepsilon.$$

20

The second last inequality follows using point (2) and (3), while the last follows using $\varepsilon \leqslant \frac{1}{5k}$. This contradicts strategy i giving at least best payoff minus ε, i.e., contradiction to (x, y) being ε-approximate well-supported equilibrium. □

Using Lemmas 5.4 and 5.5 next we show that (R, C) gives an SoSHard game. Note that $R, C \in [-1, k]$.

We rely on a key result that is a variant of the one developed by Tulsiani [Tul09]. We include an elementary proof in the appendix.

Fact 5.6 ([Tul09]). *Let* $\Psi_1(x), \Psi_2(y) = \{g_1 = 0, g_2 = 0, \dots, g_q = 0, h_1 \geqslant 0, h_2 \geqslant 0, \dots, h_r \geqslant 0\}$ *be two systems of polynomial constraints in* $x \in \mathbb{R}^{n_1}$ *and* $y \in \mathbb{R}^{n_2}$, *respectively.*

Suppose there are polynomials $f_1, f_2, \dots, f_{n_2}$ *of degree at most* t, *and sum-of-squares polynomial* $Q_1, Q_2, \dots, Q_r$ *such that for every* x *that is feasible for* Ψ_1, $y \in \mathbb{R}^{n_2}$ *defined by* $y_j = f_j(x)$ *for every* $j \in [n_2]$ *satisfies:*

1. $g_j(y) = 0$ *for every* $1 \leqslant j \leqslant q$, *and*

2. $h_j(y) = Q_j(y)$.

Then, if there exists a degree d *pseudo-expectation satisfying the constraints of* Ψ_1, *then there exists a degree* $\lfloor d/t \rfloor$-*pseudo-distribution satisfying the constraints of* Ψ_2.

Lemma 5.7. *Given the combinatorial problem k-IS on graph* $G = (V, E)$ *on* n *vertices, construct* (R, C) *as per (5.1). Game* $\frac{1}{k}(R, C)$ *is a degree* $\Omega(n)$, (ε, δ)-*SoSHard, where* $\varepsilon = \frac{1}{5k^2}$ *and* $\delta = (\frac{1}{k} + \frac{1}{2k^2})$.

Proof. First note that all payoff entries of game $(\tilde{R}, \tilde{C}) = \frac{1}{k}(R, C)$ are in $[-1, 1]$. For the *completeness* property of an SoSHard game, note that solution to independent-set gives a Nash equilibrium of $(\tilde{R}, \tilde{C})$ with payoff at least $(\frac{1}{k} + \frac{\gamma}{k^2}) = \delta$ (Lemma 5.4). Thus, pseudo-distribution for the independent set can be mapped to one for Nash equilibrium using Fact 5.6 on reductions within SoS framework. Furthermore, if the size of the independent set is fixed, then this map is low degree (Fact 5.6). Finally, the known degree $\Omega(n)$ SoS hardness for independent set of fixed $O(n)$ size [Tul09], gives degree $\Omega(n)$ pseudo-equilibrium for game $(\tilde{R}, \tilde{C})$ with payoff δ.

The *soundness* property follows using Lemma 5.5 since when both players play their last strategy, they both get payoff $\frac{1}{k} < (\delta - \varepsilon)$. □

Lemma 5.7, the fact that $k \leqslant n$, together with the SoS hardness of independent set and Fact 5.6 give Lemma 5.2.

5.2 EnumHard Games Construction

In this section we prove Lemma 5.3. For this we construct an $n \times n$ EnumHard game for a given subset $S \subset [n]$. Let $|S| = m$. We will first construct an $m \times m$ game whose $O(1/m^2)$-WNE set is near the uniform strategy.

Generalized Matching Pennies [DGP09]. This game was introduced in [DGP09], however, a characterization of its approximate equilibria was not known to the best of our knowledge. Consider a zero-sum two-player game with m strategies, where the row player wants to match with the column player and the column player does not want to match. That is if both plays $i \in [m]$ then the row player one gets payoff 1 and the column player gets -1. Otherwise, both gets zero. If $I(m)$

denotes the identity matrix of size m then, the game essentially is $(I(m), -I(m))$. It is easy to see that both players playing uniform random strategy is the only Nash equilibrium [DGP09]. Next we show that all ε-WNE are also near this unique NE.

Lemma 5.8. *For $0 \leqslant \varepsilon \leqslant \frac{1}{m}$ if (x, y) is an ε-WNE of the generalized matching-pennies game $(I(m), -I(m))$ then $\|x - \mathcal{U}_m\|_1 \leqslant m\varepsilon$ and $\|y - \mathcal{U}_m\|_1 \leqslant m\varepsilon$.*

Proof. We will argue for x, and similar argument follows for y. It suffices to show that $\|x - \mathcal{U}_m\|_\infty \leqslant \varepsilon$. To the contrary suppose $\exists i \in [m]$ such that $x_i > 1/m + \varepsilon$. Then $\exists i' \in [m]$ such that $x_{i'} < 1/m$. In that case, $i \notin \varepsilon\text{-BR}_2(x)$ and therefore $y_i = 0$. This would imply existence of a $j \in [m]$ such that $y_j \geqslant \frac{1}{(m-1)} > \varepsilon$. Now note that row player's payoff from strategy j is $y_j > \varepsilon$ while from strategy i it is zero. Therefore $i \notin \varepsilon\text{-BR}_1(y)$, a contradiction to (x, y) being an ε-WNE. $\qquad\square$

For any given $S \subseteq [n]$ with $|S| = m$, consider the following $n \times n$ game (R_S, C_S).

Dummy strategies: $\forall i \notin S, \forall j \in [n]$ set $R_S(i, j) = C_S(j, i) = -1$.
Matching pennies: $\forall i, j \in S$, set $R_S(i, j) = I(m)(i, j) + 1$ and $C_S(i, j) = -I(m)(i, j) + 3$.
Rest: $\forall i \in S$, $\forall j \notin S$ set $R_S(i, j) = 1$ and $C_S(j, i) = 2$.
Normalization: Divide R_S by 2, and C_S by 3. $\qquad(5.2)$

Lemma 5.9. *For any given $S \subseteq [n]$, game (R_S, C_S) of (5.2) is (ε, τ, c)-EnumHard for any $\varepsilon \leqslant \frac{1}{3n}$, $\tau = 3n$, and $c = 1$.*

Proof. By construction payoffs in (R_S, C_S) are in $[-1, 1]$, and both players get payoff at least $1/2$ from any strategy $i \in S$ no matter what the opponent is playing. Therefore, property (1) of EnumHard game (Definition 4.3) is satisfied.

Let $|S| = m \leqslant n$. For property (2), note that no matter what the opponent is playing, payoff from strategy $i \notin S$ is at most $-1/3$ which is way less than the payoff of $1/2$ from any strategy $i \in S$. Therefore, if (x, y) is an ε-WNE, then $\text{supp}(x), \text{supp}(y) \subseteq S$. Therefore, strategy profile (x', y'), the projection of (x, y) on to coordinates of set S, gives a 3ε-WNE of the matching pennies game $(I(m), -I(m))$. Thereby using Lemma 5.8 it follows that $\|x - \mathcal{U}_{n,S}\|_1 = \|x' - \mathcal{U}_m\|_1 \leqslant 3\varepsilon m \leqslant 3\varepsilon n$. Since $\|x - \mathcal{U}_{n,S}\|_1 = \|x' - \mathcal{U}_m\|_1 \leqslant 3\varepsilon n$, and this upper bound is $\tau\varepsilon$ in property (2), $\tau = 3n$ suffices. Similarly we can show $\|y - \mathcal{U}_{n,S}\|_1 \leqslant \tau\varepsilon$. $\qquad\square$

Finally to enable construction of a family of EnumHard games, we will need the following lemma.

Lemma 5.10. *For any two sets $S, S' \subseteq [n]$, $S \neq S'$, $\|\mathcal{U}_{n,S} - \mathcal{U}_{n,S'}\|_1 \geqslant \frac{1}{n}$.*

Proof. Without loss of generality suppose $|S| \geqslant |S'|$, then $\exists i \in S \setminus S'$. Therefore, $\|\mathcal{U}_{n,S} - \mathcal{U}_{n,S'}\|_1 \geqslant \mathcal{U}_{n,S}(i) \geqslant \frac{1}{|S|} \geqslant \frac{1}{n}$. $\qquad\square$

Now we are ready to prove Lemma 5.3 using Lemmas 5.9 and 5.10.

Proof of Lemma 5.3. We will construct a family of $n \times n$ EnumHard games. For every subset $S \subseteq [n]$, construct an $n \times n$ game (R_S, C_S) as per (5.2). Lemma 5.9 implies that (R_S, C_S) is an (ε, τ, c)-EnumHard game for $\varepsilon \leqslant \frac{1}{60n^2}$, $\tau = 3n$, and $c = 1$. Furthermore, for any two distinct sets $S, S' \subseteq [n]$, Lemma 5.10 implies $\|\mathcal{U}_{n,S} - \mathcal{U}_{n,S'}\|_1 \geqslant \frac{1}{n} > \frac{18}{c}(\tau + 1)\varepsilon$. $\qquad\square$

6 Quasipolynomial Hardness for Finding Approximate Equilibria

The goal of this section is to prove Theorem 2.18. We restate it for convenience.

Theorem 6.1 (Quasi-polynomial Hardness for Approximate NE). *For every n large enough, there's a family of $\Gamma = n^{\Omega(\log(n))}$ two-player games $\{(R_i, C_i)\}_{i=1}^{\Gamma}$ with n pure strategies for both players and all payoffs in $[-1, 1]$ such that:*

1. ***Completeness:*** *All $G_i s$ have a common degree $\Theta(\log(n))$, pseudo-equilibrium.*

2. ***Soundness:*** *For a fixed $\varepsilon = O(1)$, for any i, j and any pair of ε-NE in G_i, G_j, say (x, y) and (x', y') respectively, $x \neq x'$ and $y \neq y'$.*

The proof follows immediately from the following two constructions of hard games along Theorem 4.8.

Lemma 6.2 (Games Hard for SoS). *There exists a degree $\Omega(\log(n))$, $(\varepsilon, 1)$-**SoSHard** game (R, C) for $\varepsilon = O(1)$ and $R, C \in [-1, 1]^{n \times n}$.*

Lemma 6.3 (Games Hard for Enumeration). *For $K = O(n)$, $\varepsilon = O(1)$, $\tau = O(1)$ and $c = 1/2$, there exists a family of subsets $\mathcal{F} \subset 2^{\{0,1\}^K}$ of size at least $K^{\Omega(\log(K))}$ such that*

- *for each $S \in \mathcal{F}$ there is an (ε, τ, c)-**EnumHard** game (R_S, C_S) for $R_S, C_S \in [-1, 1]^{K \times K}$.*

- *for any pair of distinct $S, S' \in \mathcal{F}$, $\|\mathcal{U}_{cK,S} - \mathcal{U}_{cK,S'}\| > \frac{18}{c}(\tau + 1)\varepsilon$.*

6.1 Constructing an **SoSHard** game

We will next show construction of SoSHard game for a constant $\varepsilon > 0$. In [DFS16], the authors give a reduction from 3SAT to finding constant-approximate Nash equilibrium maximizing social welfare up to some additive slack. This reduction maps a n variable, $O(n)$-clause 3SAT instance into a $2^{O(\sqrt{n})}$ size two-player game. This gives a quasi-polynomial lower bound for the latter problem, assuming that n-variable, $O(n)$-clause 3SAT instances require exponential time to solve.

We simply observe that their reduction is low-degree, that is, the map that takes any satisfying assignment of the starting 3SAT instance into an approximate equilibrium maximizing social welfare is a polynomial of degree $O(\sqrt{n})$. Combined with Fact 5.6 as in the previous section, we immediately obtain the SoSHard game as required. Our reduction will in fact yield hardness of finding the weaker solution concept of ε-WNE (thus implying the hardness of ε-NE).

Lemma 6.4 (SoS Hardness for Finding Payoff Maximizing ε-NE). *There exists a game (R, C) for $R, C \in [-1, 1]^{N \times N}$ such that*

1. ***Completeness:*** *There's a degree $\Omega(\log(N))$, pseudo-equilibrium for (R, C) that has payoffs for both players at least 1.*

2. ***Soundness:*** *For every ε-NE (x, y) of (R, C), both row and column players have an individual payoff of at most $(1 - \varepsilon)$ for $\varepsilon < 1/1200$.*

For completeness, we sketch the argument of Deligkas-Fearnley-Savani [DFS16] next. The starting point is the hardness of approximation for the 3SAT problem. To show ETH-hardness, one has to start from a (worst-case hard) instance of 3SAT and applies the PCP theorem to get

23

a hard-to-approximate instance. For SoS-hardness, we only need such a result to hold for the SoS algorithm. As a result, we can rely on Grigoriev's theorem to get exponential SoS-hardness of approximating 3SAT. One can think of this as saying that ETH is unconditionally true when restricted to the SoS framework and, in fact, the hard instance is simply a random one!

Fact 6.5 (Grigoriev's Theorem). *Let ϕ be a random instance of 3SAT with $m = \Theta(n/\varepsilon^2)$ constraints and n variables. Then, with probability at least $1 - 1/n$,*

- *For every $x \in \{0,1\}^n$, the fraction of clauses satisfied in ϕ by x is at most $7/8 + \varepsilon$, and*

- *there exists a degree $\Theta(n)$ pseudo-distribution over the satisfying assignments for the instance ϕ.*

Given a hard instance of 3SAT, one transforms it into a 2 prover, *free* game.

Recall that a two prover game is defined by a distribution $\mathcal{D}$ over a finite set of question pairs (Q_1, Q_2) with a collection of admissible answer pairs and two players Alice and Bob. In a run of the game, the referee draws a question pair (Q_1, Q_2) from $\mathcal{D}$ and sends Q_1 to Alice and Q_2 to Bob and receives answers from them and accepts if the answer pair is in the list of admissible answer pairs for (Q_1, Q_2). A two prover game is said to be *free* if $\mathcal{D}$ is a product distribution.

A *strategy* for a two prover game is a pair of functions f_1 and f_2 that maps questions to an answer. The value of the strategy is the probability that $(f_1(Q_1), f_2(Q_2))$ belongs to an admissible answer pair for the question pair (Q_1, Q_2) picked at random from $\mathcal{D}$.

Definition 6.6 (Free Game from 3SAT). Let ϕ be a 3SAT instance from Fact 6.5 with m clauses and n variables. Let the clauses and variables appearing in ϕ each be partitioned into $\sqrt{n}$ different groups, say, $S_1, S_2, \ldots, S_{\sqrt{m}}$ and $T_1, T_2, \ldots, T_{\sqrt{n}}$, respectively. Then, ϕ defines a free game as follows:

Let $\mathcal{D}$ be the uniform distribution on (S_i, T_j) for $i \sim [\sqrt{m}], j \sim [\sqrt{n}]$ define the distribution on question pairs. For (S_i, T_j), an admissible answer pair is an assignment to the variables appearing in each clause in S_i and each variable in T_j such that 1) every clause in S_i is satisfied by the assignment chosen 2) for every variable that appears in some clause in S_i and in T_j, the assignments given by the answer pairs match.

Note that the number of strategies for Merlin 1 are at most $N_1 = 2^{3\sqrt{m}}$ and that for Merlin 2, at most $N_2 = 2^{\sqrt{n}}$.

The free game $\mathcal{F}$ can be thought of as a cooperative two player game as follows. The row player has as their strategies the set of all ordered pairs of questions to Merlin 1 and admissible answers for that question. Similarly, the column player has as strategies, all possible questions to Merlin 2 and admissible answers for those questions. Finally, a pair of ordered pairs of questions and answers receive a payoff of 1 (for both row and column player) iff the verifier accepts the two answers on those two questions and 0 otherwise. Observe that the size of the cooperative game produced from a free game is at most $\sqrt{m}N_1 \times \sqrt{n}N_2$.

The following completeness result holds for every free game produced from 3SAT instances from Fact 6.5.

Lemma 6.7 (SoS Completeness for Free Game). *Let ϕ be the 3SAT instance with $m = \Theta(n/\varepsilon^2)$ clauses and n variables from Fact 6.5. Let $\mathcal{F}$ be the free game constructed in Definition 6.6 from ϕ. Then, there's a pseudo-equilibrium of degree $\Theta(n)$ for the cooperative game defined by $\mathcal{F}$ with payoff for each player $= 1$.*

Proof. We will start from the pseudo-distribution of degree $\Theta(n)$ given to us by Fact 6.5 that satisfies every constraint of ϕ and use Fact 5.6. To do this, we need to build a low-degree map from any

24

satisfying assignment x of ϕ to a strategy for the Merlins in the free game $\mathcal{F}$. This is very simple: for any question set (S_i, T_j) - the strategy is described by a bit string of length $3\sqrt{m}$ and $\sqrt{n}$ respectively - where the bit string corresponds to the assignment for the variables appearing in the clauses in S_i and variables in T_j respectively. For the strategy for the cooperative game we use the following mixed strategy: Merlin 1 chooses a question from X uniformly at random and then chooses the answer corresponding to the assignment x for that question. Similarly, Merlin 2 chooses a question from Y uniformly at random and then chooses the answer corresponding to x for that questions. Using x_i as the assignment for the variable i wherever it appears gives a degree 1 map into strategies of $\mathcal{F}$.

It is easy to see that since x (purportedly) satisfies all constraints of ϕ, it satisfies all clauses in S_i and further, the assignments to variables appearing in both S_i and T_j match. Thus, we obtain a degree 1 map from any satisfying assignment x into a strategy for $\mathcal{F}$ with value 1. Further, it is easy to see (since the maximum possible payoff in the game is 1) that the strategy constructed above satisfies all constraints of being an NE. Thus, we have exhibited a degree $O(\sqrt{n})$ map from any satisfying assignment x to a NE of the cooperative game associated with $\mathcal{F}$. By Fact 5.6, we immediately obtain a pseudo-distribution of degree $\Theta(n)$ on strategies with value 1 for $\mathcal{F}$.

$\square$

The following soundness result holds for the partitions $S_1, S_2, \ldots, S_{\sqrt{m}}$ and $T_1, T_2, \ldots, T_{\sqrt{n}}$ chosen appropriately (we invite the reader to find the details of the reduction in Deligkas-Fearnley-Savani [DFS16]):

Fact 6.8 (Soundness, [DFS16]). *There is no strategy for $\mathcal{F}$ produced from ϕ in Fact 6.5 with value exceeding $1 - \Omega(1)$.*

Deligkas-Fearnley-Savani give a construction of a game (R', C') by taking the cooperative game (R, C) produced from the free game as above and using it to fill the first $\sqrt{m}N_1 \times \sqrt{n}N_2$ block and then adding H rows and columns for H being $2^{|Y|} \approx 2^{\sqrt{n}}$. The off diagonal blocks are filled with a carefully constructed zero sum game and its transpose. Finally, the bottom diagonal $H \times H$ block is filled with the all zero matrix. It is immediate from the simple argument above that the pseudo-equilibrium described above for $\mathcal{F}$ is in fact a ε-approximate pseudo-equilibrium in (R', C') - the only required fact in this argument is that the equilibrium strategy $\mathbf{y}$ ($\mathbf{x}$ respectively) satisfy $\sum_b \mathbf{y}(y, b) = 1$ for every question y to Merlin 2. This is a linear polynomial equality constraint that carries over to the pseudo-distribution in a straightforward way.

We omit the simple and standard argument here and summarize the result that immediately follows below:

Fact 6.9 (Corollary of [DFS16]). *Let $\mathcal{F}$ be a free game with $\sqrt{m}$ ($\sqrt{n}$, respectively) possible questions and N_1 (N_2, respectively) possible strategies for Merlin 1 (2, respectively.) Then, for $\varepsilon = O(1)$ there's a two player game with $\sqrt{m} \cdot N_1$ strategies for the row player and $\sqrt{n} \cdot N_2$ strategies for the column payer such that:*

1. **Completeness:** *There's a ε-approximate pseudo-equilibrium of degree $\Theta(n) = \Omega(\log(N_1 + N_2))$ where both players receive a payoff of at least 1.*

2. **Soundness:** *For every ε-approximate NE of the game, the payoff of either player is at most $(1 - \varepsilon)$ for $\varepsilon < 1/1200$.*

6.2 Constructing a family of **EnumHard** games

In this section we prove Lemma 6.3 by constructing a family of EnumHard games for a constant ε. Essentially, we will construct a family of two-player games, each parameterized by a subset S of strategies such that: (i) both players' strategies at Nash equilibrium is largely supported on the set S, and (ii) for a constant ε, all its ε-WNE are near the actual NE. Daskalakis and Papadimitriou [DP09] gave a construction that ensures row-player's all WNE strategies are near the uniform distribution over the subset. Constructing such a game for $n^{log(n)}$ carefully selected subsets, they showed non-existence of an oblivious PTAS for WNE that uses only Verification Oracle 2.9, i.e., checks if given strategy profile is a WNE . However, in their construction the column player's NE (WNE) strategies are common across games. Therefore it breaks under our more powerful Verification Oracle 2.12, *i.e.,* given any *one* player's strategy the oracle checks if there exists a ε-NE corresponding to the given strategy.

Since we need hardness against Oracle 2.12, construction of [DP09] can not be used directly. We will have to ensure that, for each player, her NE strategy sets are far apart across games. Our construction builds on that of [DP09], and therefore first we discuss the latter and its result.

Construction of [DP09]. Let l be an even integer and $n = \binom{l}{l/2}$. For each $S \subset [n]$ of size l, we will construct an $n \times n$ game. Fix such a subset S, and denote the corresponding game by (A_S, B_S). In this game, row player cares about only strategies from S, *i.e.,* strategies outside of S are strictly dominated by that of S. Since $|S|$ is l, number of $(l/2)$ sized subsets of S are exactly n. The n strategies of the column player correspond to these $(l/2)$ sized subsets of set S. In other words, if $S_1, \ldots, S_{\binom{l}{l/2}}$ are all the $(l/2)$ subsets of S, then think of column player's strategies being indexed by these subsets. The payoffs in column j, corresponding to subset S_j of S, are:

- For each $i \notin S$, set $A_S(i, j) = -1$ and $B_S(i, j) = 1$.

- For each $i \in S \cap S_j$, set $A_S(i, j) = 1$ and $B_S(i, j) = 0$.

- For each $i \in S \setminus S_j$, set $A_S(i, j) = 0$ and $B_S(i, j) = 1$.

Theorem 6.10. *[DP09]*

(a) *For any given $\varepsilon < 1$, if (x, y) is an ε-WNE of game (A_S, B_S), then (i) for all $i \notin S$, $x_i = 0$, and (ii) $\|x - \mathcal{U}_{n,S}\|_1 \leqslant 8\varepsilon$.*

(b) *For $V = \{\mathcal{U}_{n,S} \mid S \subseteq [n], |S| = l\}$ we have $|V| = \Omega(n^{0.8 \log_2 n})$. And for any $\gamma > 0$ there exists $V' \subset V$ of size $\Omega(n^{(0.8 - 2\gamma\varepsilon) \log_2 n})$ such that for any two vectors $\mathcal{U}_{n,S}, \mathcal{U}_{n,S'} \in V'$ we have $\|\mathcal{U}_{n,S'} - \mathcal{U}_{n,S}\|_1 > \gamma\varepsilon$.*

Note that, column player playing each of its n strategies uniformly at random is her NE strategy in game (A_S, B_S) for all $S \subset [n]$ of size l. Thus querying Oracle 2.12 with strategy $y = \mathcal{U}_n$ for the column player will return a NE of (A_S, B_S). To prevent this next we extend this construction that uses the above game combined with another copy of it where roles of the two players are switched. For the construction to work as desired, we need both players to put non-trivial probability mass on both games, and to achieve this we augment this game with a matching-pennies game type construction.

Our Construction. Again fix an even integer l and let $n = \binom{l}{l/2}$. Fix an l sized subset $S \subset [n]$, and corresponding game (A_S, B_S) as described above. We construct a $2n \times 2n$ game (R_S, C_S) as follows: if we think of matrices of this game as 2×2 block matrix, each block of size $n \times n$, then

$$R_S = \begin{bmatrix} 2 + A_S & -2 \\ -2 & 2 + B_S^\top \end{bmatrix} \qquad C_S = \begin{bmatrix} -2 + B_S & 2 \\ 2 & -2 + A_S^\top \end{bmatrix} \qquad (6.1)$$

In (6.1), the constants 2 and -2 represents block matrices of size $n \times n$ with all entries set to 2 and -2 respectively. For any strategy profile (x, y) of game (R_S, C_S), let us divide each of x and y into two vectors corresponding to first n and last n strategies. These are $x^L = (x_1, \ldots, x_n)$, $x^R = (x_{(n+1)}, \ldots, x_{2n})$, $y^L = (y_1, \ldots, y_n)$, and $y^R = (y_{(n+1)}, \ldots, y_{2n})$. For any vector $\mathbf{v}$, we will use $\sigma(\mathbf{v})$ to denote sum of its coordinates $(\sum_i v_i)$.

Using the matching pennies property of game (R_S, C_S), next we show that both x^L and x^R have enough probability mass, and similarly y^L and y^R.

Lemma 6.11. *For an $\varepsilon \in [0, 1)$, let (x, y) be an ε-WNE of game (R_S, C_S). Then, $\frac{9}{19} - \frac{\varepsilon}{9} \leqslant \sigma(x^L), \sigma(x^R), \sigma(y^L), \sigma(y^R) \leqslant \frac{10}{19} + \frac{\varepsilon}{9}$.*

Proof. First we show that $\sigma(x^L), \sigma(x^R), \sigma(y^L), \sigma(y^R) > 0$. To the contrary suppose $\sigma(x^R) = 0$, then column player gets payoff at most -1 from any strategy $j \leqslant n$, while gets payoff of 2 from any $k > n$. Therefore $\sigma(y^L) = 0$. In that case, the row player gets payoff at most -1 from any $i \leqslant n$, while gets at least 2 payoff from any $h > n$. Since there exists $i \leqslant n$ such that $x_i > 0$, this contradicts (x, y) being ε-WNE. Similar argument follows for each of $\sigma(x^L), \sigma(y^L), \sigma(y^R)$ being zero.

Next we will argue that $\frac{9}{19} - \frac{\varepsilon}{9} \leqslant \sigma(y^L) \leqslant \frac{10}{19} + \frac{\varepsilon}{9}$, and then using $\sigma(y^R) = 1 - \sigma(y^L)$ the inequalities will follow for $\sigma(y^R)$ as well. For this we need properties of payoffs from A_S and B_S blocks of R_S.

Claim 6.12. For all $k \leqslant n$ we have $e_k^\top A_S y^L \leqslant 1$ and $e_k^\top B_S^\top y^R \leqslant 1$, and there exists an $i, i' \leqslant n$ such that $e_i^\top A_S y^L \geqslant \frac{\sigma(y^L)}{2}$ and $e_{i'}^\top B_S^\top y^R \geqslant \frac{\sigma(y^R)}{2}$.

Proof. The first part follows by the fact that coordinates of both A_S and B_S are at most one. For the second part, note that for each column $j \in [n]$ of A_S exactly $(l/2)$ coordinates are 1 out of l coordinates of S. Therefore, summing up payoffs corresponding to strategies of S, we have

$$\sum_{i \in S} e_i^\top A_S y^L = \sum_{j \leqslant [n]} y^L * (l/2) = \sigma(y^L)(l/2)$$

Since $|S|$ is l, among the strategies of S there is at least one $i \in S$ such that $e_i^\top A_S y^L \geqslant \frac{\sigma(y^L)}{2}$. For the second part, a row of $B_S^\top$ is a column of B_S. Note that every entry in B_S is at least zero, and in a column all entries corresponding to rows in S are one. Let us divide y^R further into coordinates corresponding to set S, namely $y^{R,S}$, and those outside S, namely $y^{R,\bar{S}}$. In block B_S the role of first player is that of column-player in game (A_S, B_S). Therefore, strategies $(n+1), \ldots, 2n$ correspond to the $n = \binom{l}{(l/2)}$ subsets of S. There is one such set that correspond to largest $(l/2)$ coordinates of $y^{R,S}$. Let this be $S_{i'}$. Then, $e_{i'}^\top B_S^\top y^R \geqslant \sigma(y^{R,\bar{S}}) + \sigma(y^{R,S})/2 \geqslant \sigma(y^R)/2$. $\qquad\square$

Consider $i, i' \leqslant n$ as in Claim 6.12. Since $\sigma(x^L), \sigma(x^R) > 0$, there exists $h, h' \leqslant n$ such that $x_h, x_{(n+h')} > 0$. Row player's payoffs from strategies $i, (n + i'), h$, and $(n + h')$ are:

$$e_h^\top R_S y = 2\sigma(y^L) + e_h^\top A_S y^L - 2\sigma(y^R), \quad e_{(n+i')}^\top R_S y = -2\sigma(y^L) + e_{i'}^\top B_S^\top y^R + 2\sigma(y^R)$$
$$e_i^\top R_S y = 2\sigma(y^L) + e_i^\top A_S y^L - 2\sigma(y^R), \quad e_{(n+h')}^\top R_S y = -2\sigma(y^L) + e_{h'}^\top B_S^\top y^R + 2\sigma(y^R)$$

Since (x, y) is an ε-WNE we must have $e_h^\top R_S y \geqslant e_{(n+i')}^\top R_S y - \varepsilon$ and $e_{(n+h')}^\top R_S y \geqslant e_i^\top R_S y - \varepsilon$. Using Claim 6.12, the former gives,

$$3\sigma(y^L) - 2\sigma(y^R) \geqslant -2\sigma(y^L) + \frac{\sigma(y^R)}{2} + 2\sigma(y^R) - \varepsilon \;\; \Rightarrow \;\; 5\sigma(y^L) \geqslant \frac{9}{2}\sigma(y^R) - \varepsilon$$
$$\Rightarrow \;\; \sigma(y^L) \geqslant \frac{9}{10}\sigma(y^R) - \frac{\varepsilon}{5} = \frac{9}{10}(1 - \sigma(y^L)) - \frac{\varepsilon}{5}$$
$$\Rightarrow \;\; (1 + \frac{9}{10})\sigma(y^L) \geqslant \frac{9}{10} - \frac{\varepsilon}{5}$$
$$\Rightarrow \;\; \sigma(y^L) \geqslant \frac{9}{19} - \frac{2}{19}\varepsilon \geqslant \frac{9}{19} - \frac{\varepsilon}{9}$$

Similarly, using $e_{(n+h')}^\top R_S y \geqslant e_i^\top R_S y - \varepsilon$ together with Claim 6.12 we will get $\sigma(y^R) \geqslant \frac{9}{10}\sigma(y^L) - \frac{\varepsilon}{5}$. Simplifying this using $\sigma(y^R) = 1 - \sigma(y^L)$ gives $\sigma(y^L) \leqslant \frac{10}{19} + \frac{\varepsilon}{9}$.

Bounds on $\sigma(x^L)$ and $\sigma(x^R)$ can be obtained through similar arguments for the payoff of the column-player w.r.t. C_S. □

For any $\varepsilon < 1$ the bounds obtained in Lemma 6.11 implies lower bound of at least $\frac{1}{3}$ and upper bound of at most $\frac{2}{3}$ on the each of $\sigma(x^L)$, $\sigma(x^R)$, $\sigma(y^L)$, and $\sigma(y^R)$.

Lemma 6.13. *For an $\varepsilon \in [0, 1/3)$, let (x, y) be an ε-WNE of game (R_S, C_S). Then, $(\frac{x^L}{\sigma(x^L)}, \frac{y^L}{\sigma(y^L)})$ and $(\frac{y^R}{\sigma(y^R)}, \frac{x^R}{\sigma(x^R)})$ are 3ε-WNE of game (A_S, B_S).*

Proof. If x_i for an $i \leqslant n$, then it must be the case that $e_i^\top R_S y \geqslant e_k^\top R_S y - \varepsilon$ for all $k \leqslant n$. This implies,

$$e_i^\top A_S y^L \geqslant e_k^\top A_S y^L - \varepsilon \Rightarrow \frac{e_i^\top A_S y^L}{\sigma(y^L)} \geqslant \frac{e_k^\top A_S y^L}{\sigma(y^L)} - \frac{\varepsilon}{\sigma(y^L)} \Rightarrow (e_i^\top A_S \frac{y^L}{\sigma(y^L)}) \geqslant (e_k^\top A_S \frac{y^L}{\sigma(y^L)}) - 3\varepsilon$$

The last inequality follows using Lemma 6.11. Similarly, we can argue that if $y_j > 0$ for a $j \leqslant n$, then for all $k \leqslant n$ we have $(\frac{x^{L\top}}{\sigma(x^L)} B_S e_j) \geqslant (\frac{x^{L\top}}{\sigma(x^L)} B_S e_k) - 3\varepsilon$. These two together implies $(\frac{x^L}{\sigma(x^L)}, \frac{y^L}{\sigma(y^L)})$ is a 3ε-WNE of game (A_S, B_S).

In the lower diagonal block, the game is (B_S^T, A_S^T), or in other words the roles of the two players are switched. Therefore, it follows that if $x_i > 0$ for an $i > n$ then for all $k > n$ we have $(B_S^T \frac{y^R}{\sigma(y^R)} e_i) = (\frac{y^{R\top}}{\sigma(y^R)} B_S e_i) \geqslant (\frac{y^{R\top}}{\sigma(y^R)} B_S e_k) - 3\varepsilon$. And similarly, if $y_j > 0$ for $j > n$ then for all $k > n$ we have $(e_j^\top A_S \frac{x^R}{\sigma(x^R)}) \geqslant (e_k^\top A_S \frac{x^R}{\sigma(x^R)}) - 3\varepsilon$. These two together implies $(\frac{y^R}{\sigma(y^R)}, \frac{x^R}{\sigma(x^R)})$ are 3ε-WNE of game (A_S, B_S). The lemma follows. □

The above lemma together with Theorem 6.10 implies that for any ε-WNE (x, y), vector x^L is near uniform distribution over strategies of S scaled by $\sigma(x^L)$. In order to show non-overlapping equilibrium sets across games corresponding to different S, we will need that x^L is near such a uniform distribution scaled by $\frac{1}{2}$. For this we need a tighter upper-lower bound on $\sigma(x^L)$, that we show next. The main missing component here is a tighter upper bound on the payoffs from (A_S, B_S) proved in the next lemma.

Lemma 6.14. *For an $\varepsilon \in [0, 1/3)$, if (x, y) is an ε-WNE of game (R_S, C_S), then $\forall i \leqslant n$, $(e_i^\top A_S y^L) \leqslant \frac{\sigma(y^L)}{2} + 4\varepsilon$ and $(x^{L\top} B_S e_i) \leqslant \frac{\sigma(x^L)}{2} + 4\varepsilon$. And similarly, $(e_i^\top B_S^T y^R) \leqslant \frac{\sigma(y^L)}{2} + 4\varepsilon$ and $(e_i^\top A_S x^R) \leqslant \frac{\sigma(x^R)}{2} + 4\varepsilon$ for each $i \in [n]$.*

Proof. Let $\tilde{x} = \frac{x^L}{\sigma(x^L)}$ and $\tilde{y} = \frac{y^L}{\sigma(y^L)}$. Lemma 6.13 implies that $(\tilde{x}, \tilde{y})$ is a 3ε-WNE of game (A_S, B_S). Then due to Theorem 6.10 part (a), we have $\tilde{x}_i = 0$ for all $i \notin S$. Since sum of the payoffs for rows in S is always 1 in this game, we have $\tilde{x}^\top \top (A_S + B_S)\tilde{y} = 1$. There are two possible scenarios:

Case I - $\tilde{x}^\top A_S \tilde{y} \leqslant \frac{1}{2}$. Then it must be the case that $\forall i \leqslant n$, $(e_i^\top A_S \tilde{y}) \leqslant \frac{1}{2} + 3\varepsilon$, or else $(\tilde{x}, \tilde{y})$ is not a 3ε-WNE of (A_S, B_S). Multiplying by $\sigma(y^L)$ gives, $(e_i^\top A_S y^L) \leqslant \frac{\sigma(y^L)}{2} + \sigma(y^L)3\varepsilon \leqslant \frac{\sigma(y^L)}{2} + 2\varepsilon$ for each $i \in [n]$. The last inequality uses $\sigma(y^L) \leqslant \frac{2}{3}$ from Lemma 6.11.

To get similar bound on $(x^{L^\top} B_S e_i)$, note that Claim 6.12 gives existence of $i \in [n]$ such that $(e_i^\top A_S \tilde{y}) \geqslant \frac{1}{2}$. Therefore $\tilde{x}^\top A_S \tilde{y} \geqslant \frac{1}{2} - 3\varepsilon$, implying $\tilde{x}^\top B_S \tilde{y} \leqslant \frac{1}{2} + 3\varepsilon$. Therefore, again 3ε-WNE condition gives that for each $i \in [n]$ we should have $(\tilde{x}^\top B_S e_i) \leqslant \frac{1}{2} + 6\varepsilon$. Through similar analysis as above we would get that for each $i \in [n]$, $(x^{L^\top} B_S e_i) \leqslant \frac{\sigma(x^L)}{2} + 4\varepsilon$.

Case II - $\tilde{x}^\top B_S \tilde{y} \leqslant \frac{1}{2}$. Again by similar analysis as Case I, we will get that for each $i \in [n]$ we have $(e_i^\top A_S y^L) \leqslant \frac{\sigma(y^L)}{+} 4\varepsilon$ and $(x^{L^\top} B_S e_i) \leqslant \frac{\sigma(x^L)}{2} + 2\varepsilon$.

The second part follows similarly using the fact that $(\frac{y^R}{\sigma(y^R)}, \frac{x^R}{\sigma(x^R)})$ is also a 3ε-WNE of game (A_S, B_S) from Lemma 6.13. $\qquad\square$

Now we are ready to prove a tighter lower/upper bound on $\sigma(x^L), \sigma(x^R), \sigma(y^L), \sigma(y^R)$.

Lemma 6.15. *Given an* $\varepsilon \in [0, 1)$, *if* (x, y) *is an* ε-*WNE of game* (R_S, C_S) *then* $\frac{1}{2} - \varepsilon \leqslant \sigma(x^L), \sigma(x^R), \sigma(y^L), \sigma(y^R) \leqslant \frac{1}{2} + \varepsilon$.

Proof. The proof follows similar argument as of the proof of Lemma 6.11. We will first show the claim for $\sigma(y^L)$ and thereby it will follow for $\sigma(y^R)$ as well since $y^R = 1 - y^L$. Let $i, i' \in [n]$ be the indices from Claim 6.12. Let $h, h' \in [n]$ be such that $x_h, x_{(n+h')} > 0$. Then, $(e_h^\top R_S y) \geqslant (e_{(n+i')}^\top R_S y) - \varepsilon$ gives

$$2\sigma(y^L) + (e_i^\top A_S y^L) - 2\sigma(y^R) \geqslant -2\sigma(y^L) + (e_{i'}^\top B_S^\top y^R) + 2\sigma(y^R) - \varepsilon \Rightarrow 4\sigma(y^L) + \frac{\sigma(y^L)}{2} + \varepsilon \geqslant 2\sigma(y^R) + \frac{\sigma(y^R)}{2} - \varepsilon$$

The last inequality uses $(e_i^\top A_S y^L) \leqslant \frac{\sigma(y^L)}{2} + 4\varepsilon$ from Lemma 6.14. Simplifying the above gives $\sigma(y^L) \geqslant \sigma(y^R) - 2\varepsilon \Rightarrow \sigma(y^L) \geqslant \frac{1}{2} - \varepsilon$. Similarly, simplifying $(e_{(n+h')}^\top R_S y) \geqslant (e_i^\top R_S y) - \varepsilon$ by making use of $(e_i^\top A_S y^L) \geqslant \frac{\sigma(y^L)}{2}$ from Claim 6.12 and $(B_S^\top y^R) \leqslant \frac{\sigma(y^R)}{2} + 4\varepsilon$ from Lemma 6.14, we get $\sigma(y^L) \leqslant \frac{1}{2} + \varepsilon$.

Bounds for $\sigma(x^L), \sigma(x^R)$ follows using similar argument for the payoffs of the column player. $\quad\square$

Now we are ready to show that vector x^L and y^R corresponding to any ε-WNE are near uniform distribution on set S scaled by $1/2$.

Lemma 6.16. *For an* $\varepsilon \in [0, 1/3)$, *if* (x, y) *is an* ε-*WNE of game* (R_S, C_S), *then* (i) $\text{supp}(x^L), \text{supp}(y^R) \subseteq S$, *and* (ii) $\|x^L - \mathcal{U}_{n,S}/2\|_1 \leqslant 17\varepsilon$ *and* $\|y^R - \mathcal{U}_{n,S}/2\|_1 \leqslant 17\varepsilon$.

Proof. From Lemma 6.13 we have that $(x^L/\sigma(x^L), y^L/\sigma(y^L))$ is a 3ε-WNE of game (A_S, B_S). Then part (a) of Theorem 6.10 implies $\text{supp}(x^L) \subset S$ and $\|x^L/\sigma(x^L) - \mathcal{U}_{n,S}\|_1 \leqslant 24\varepsilon$. Multiplying by $\sigma(x^L)$ gives $\|x^L - \sigma(x^L)\mathcal{U}_{n,S}\|_1 \leqslant \sigma(x^L)24\varepsilon$. Using the fact that $\sigma(x^L) \geqslant \frac{1}{2} - \varepsilon$ from Lemma 6.15, we get $\|x^L - \frac{1}{2}\mathcal{U}_{n,S}\|_1 - \varepsilon \leqslant 16\varepsilon \Rightarrow \|x^L - \frac{1}{2}\mathcal{U}_{n,S}\|_1 \leqslant 17\varepsilon$.

Similarly, using $(y^R/\sigma(y^R), x^R/\sigma(x^R))$ being a 3ε-WNE of game (A_S, B_S) from Lemma 6.13, together with part (a) of Theorem 6.10 and lower bound $(1/2 - \varepsilon)$ on $\sigma(y^R)$ from Lemma 6.15, we can show that $\text{supp}(y^R) \subseteq S$ and $\|y^R - \frac{1}{2}\mathcal{U}_{n,S}\|_1 \leqslant 17\varepsilon$. $\qquad\square$

Recall that our final goal is to construct a family of EnumHard games for appropriate parameters of ε, τ and c such that their ε-WNE do not intersect for a constant ε. Note that all payoff entries of

an EnumHard game (Definition 4.3) are in $[-1, 1]$, and the first requirement is that no matter what the opponent plays, a player always has a strategy that gives at least half payoff. In order to ensure these we need to scale game (R_S, C_S) additive and multiplicatively by a constant. The third part requires renaming of the strategies of the column player so that vectors y^L and y^R get swapped.

$$\forall i \in [2n], \quad \forall j \in [n], \quad \begin{aligned} &\tilde{R}_S(i, j) = (R_S(i, n + j) + 3)/6, \quad \tilde{R}_S(i, n + j) = (R_S(i, j) + 3)/6, \\ &\tilde{C}_S(i, j) = (C_S(i, n + j) + 3)/6, \quad \tilde{C}_S(i, n + j) = (C_S(i, j) + 3)/6 \end{aligned} \tag{6.2}$$

Lemma 6.17. *For an $\varepsilon \in [0, 1/26)$, $\tau = 102$ and $c = 1/2$, game $(\tilde{R}_S, \tilde{C}_S)$ is an (ε, τ, c)-EnumHard game.*

Proof. By construction, every entry in (A_S, B_S) is either $0, 1$ or -1, and thereby from (6.1) every entry in (R_S, C_S) is in $[-3, 3]$. Thus, $\tilde{R}_S, \tilde{C}_S \in [-1, 1]$. Note that we have swapped the first n and last n strategies of the column player. This is just renaming of the strategies.

For property (1) of the EnumHard game, consider i and i' of Claim 6.12. Note that the proof of this claim does not use ε-WNE property of profile (x, y). For any given profile y of the column-player, the payoff of the row player from strategy i is:

$$(e_i^\top \tilde{R}_S y) = 1/6\sigma(y^L) + 5/6\sigma(y^R) + 1/12\sigma(y^R) \geqslant \frac{13}{12}(1/2 - \varepsilon) \geqslant 1/2$$

The last inequality follows using $\varepsilon < 1/26$. Similarly, we can show that no matter what the row-player plays, column player has a strategy that gives at least $\frac{1}{2}$ payoff.

For Property (2) let (x, y) be an ε-WNE of game $(\tilde{R}_S, \tilde{C}_S)$. Then, for $y' = (y^R, y^L)$, (x, y) is a 6ε-WNE of game (R_S, C_S). Then, part (i) of Lemma 6.16 ensures $\text{supp}(x^L), \text{supp}(y^L) \subset S$ implying Property $(2.a)$ of EnumHard game. While the second part ensures $\|x^L - \frac{1}{2}\mathcal{U}_{n,S}\|_1, \|y^L - \frac{1}{2}\mathcal{U}_{n,S}\|_1 \leqslant 17(6\varepsilon) = 102\varepsilon$ implying Property $(2.b)$. □

Constructing a family. Finally, to apply Theorem 4.8, its hypothesis (h_2) requires a family of EnumHard games parameterized by set S with certain properties. Using the EnumHard game of the previous section next we will construct such a family with $n^{O(\log(n))}$ many games. The next lemma follows simply using part (2) of Theorem 6.10.

Lemma 6.18. *Let $\mathcal{S} = \{S \subset [n] \mid |S| = l\}$ and $\varepsilon < 10^{-5}$. There exists $(n^{0.75 \log(n)})$ sized subset $\mathcal{S}'$ of $\mathcal{S}$ such that for any $S, S' \in \mathcal{S}'$, $S \neq S'$, we have $\|\mathcal{U}_{n,S} - \mathcal{U}_{n,S'}\|_1 > \frac{18}{c}(\tau + 1)\varepsilon$, where $\tau = 102$ and $c = 1/2$.*

Now we have all ingredients now to prove Lemma 6.3.

Proof of Lemma 6.3. Set $\varepsilon < 10^{-5}$, $\tau = 102$ and $c = 1/2$ like in Lemma 6.18. The lemma implies existence of a family of subsets $\mathcal{S}'$ of set $[n]$ such that $\|\mathcal{U}_{n,S} - \mathcal{U}_{n,S'}\|_1 > \frac{18}{c}(\tau + 1)\varepsilon$. For every set $S \in \mathcal{S}'$ we know that $|S| = l$. Therefore, game $(\tilde{R}_S, \tilde{C}_S)$ of (6.2) where (R_S, C_S) is from (6.1), is an (ε, τ, c)-EnumHard due to Lemma 6.17. By construction, we have $(\tilde{R}_S, \tilde{C}_S) \in [-1, 1]^{n \times n}$. □

References

[ABS15]   Sanjeev Arora, Boaz Barak, and David Steurer, *Subexponential algorithms for unique games and related problems*, J. ACM **62** (2015), no. 5, Art. 42, 25. MR 3424199 1

[AGMS11]   Bharat Adsul, Jugal Garg, Ruta Mehta, and Milind Sohoni, *Rank-1 bimatrix games: A homeomorphism and a polynomial time algorithm*, ACM Symposium on the Theory of Computing (STOC), 2011, pp. 195–204. 1

[AIM14]   Scott Aaronson, Russell Impagliazzo, and Dana Moshkovitz, *AM with multiple merlins*, IEEE Conference on Computational Complexity, IEEE Computer Society, 2014, pp. 44–55. 13

[AKV05]   Tim Abbott, Daniel Kane, and Paul Valiant, *On the complexity of two-player win-lose games*, IEEE Annual Symposium on Foundations of Computer Science, 2005, pp. 113–122. 1

[ALSV13]   Noga Alon, Troy Lee, Adi Shraibman, and Santosh Vempala, *The approximate rank of a matrix and its algorithmic applications*, STOC'13—Proceedings of the 2013 ACM Symposium on Theory of Computing, ACM, New York, 2013, pp. 675–684. MR 3210829 3

[AOW15]   Sarah R. Allen, Ryan O'Donnell, and David Witmer, *How to refute a random CSP*, 2015 IEEE 56th Annual Symposium on Foundations of Computer Science—FOCS 2015, IEEE Computer Soc., Los Alamitos, CA, 2015, pp. 689–708. MR 3473335 1

[ARV09]   Sanjeev Arora, Satish Rao, and Umesh Vazirani, *Expander flows, geometric embeddings and graph partitioning*, J. ACM **56** (2009), no. 2, Art. 5, 37. MR 2535878 1, 2

[Bar15]   Siddharth Barman, *Approximating Nash equilibria and dense bipartite subgraphs via an approximate version of Caratheodory's theorem*, ACM Symposium on the Theory of Computing (STOC), 2015, pp. 361–369. 1

[BBH+12]   Boaz Barak, Fernando G. S. L. Brandão, Aram W. Harrow, Jonathan Kelner, David Steurer, and Yuan Zhou, *Hypercontractivity, sum-of-squares proofs, and their applications [extended abstract]*, STOC'12—Proceedings of the 2012 ACM Symposium on Theory of Computing, ACM, New York, 2012, pp. 307–326. MR 2961513 1

[BBKK18]   Boaz Barak, Zvika Brakerski, Ilan Komargodski, and Pravesh K. Kothari, *Limits on low-degree pseudorandom generators (or: Sum-of-squares meets program obfuscation)*, Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II, 2018, pp. 649–679. 1

[BHK+16]   Boaz Barak, Samuel B. Hopkins, Jonathan A. Kelner, Pravesh Kothari, Ankur Moitra, and Aaron Potechin, *A nearly tight sum-of-squares lower bound for the planted clique problem*, Electronic Colloquium on Computational Complexity (ECCC) **23** (2016), 58. 1, 13

[BKS14]   Boaz Barak, Jonathan A. Kelner, and David Steurer, *Rounding sum-of-squares relaxations*, STOC, ACM, 2014, pp. 31–40. 1, 3

[BKS15]   Boaz Barak, Jonathan A. Kelner, and David Steurer, *Dictionary learning and tensor decomposition via the sum-of-squares method [extended abstract]*, STOC'15—Proceedings of the 2015 ACM Symposium on Theory of Computing, ACM, New York, 2015, pp. 143–151. MR 3388192 1

[BKS17]   Boaz Barak, Pravesh K. Kothari, and David Steurer, *Quantum entanglement, sum-of-squares and the log-rank conjecture*, Symposium on Theory of Computing, STOC (2017). 1

[BKW14]   Mark Braverman, Young Kun-Ko, and Omri Weinstein, *Approximating the best nash equilibrium in $n^{o(\log n)}$-time breaks the exponential time hypothesis*, Electronic Colloquium on Computational Complexity (ECCC) **21** (2014), 92. 13

[BKW15]   Mark Braverman, Young Kun Ko, and Omri Weinstein, *Approximating the best Nash equilibrium in $no(\log n)$-time breaks the exponential time hypothesis*, ACM-SIAM Annual Symposium on Discrete Algorithms (SODA), 2015, pp. 970–982. 2

[BM16]   Boaz Barak and Ankur Moitra, *Noisy tensor completion via the sum-of-squares hierarchy*, COLT, JMLR Workshop and Conference Proceedings, vol. 49, JMLR.org, 2016, pp. 417–445. 1

[BR17]   Yakov Babichenko and Aviad Rubinstein, *Communication complexity of approximate Nash equilibria*, ACM Symposium on the Theory of Computing (STOC), 2017. 1

[BRS11]   Boaz Barak, Prasad Raghavendra, and David Steurer, *Rounding semidefinite programming hierarchies via global correlation*, 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science—FOCS 2011, IEEE Computer Soc., Los Alamitos, CA, 2011, pp. 472–481. MR 2932723 1, 2

[BS]   Boaz Barak and David Steurer, *Proofs, beliefs, and algorithms through the lens of sum-of-squares*, Note for the course co-taught by Boaz Barak and Pablo Parrilo at Harvard-MIT, and David Steurer and Pravesh Kothari at Princeton. 4

[CDT09]   X. Chen, X. Deng, and S.-H. Teng, *Settling the complexity of computing two-player Nash equilibria*, Journal of the ACM **56(3)** (2009). 1, 8, 10, 15

[CS08]   Vincent Conitzer and Tuomas Sandholm, *New complexity results about Nash equilibria*, Games and Economic Behavior **63** (2008), no. 2, 621–641. 2

[DFS16]   Argyrios Deligkas, John Fearnley, and Rahul Savani, *Inapproximability results for approximate nash equilibria*, 12th International Conference Web and Internet Economics (WINE), 2016, pp. 29–43. 2, 10, 13, 23, 25

[DGP09]   C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou, *The complexity of computing a Nash equilibrium*, SIAM Journal on Computing **39(1)** (2009), 195–259. 1, 21, 22

[DP09]   Constantinos Daskalakis and Christos H. Papadimitriou, *On oblivious PTAS's for Nash equilibrium*, ACM Symposium on the Theory of Computing (STOC), 2009, pp. 75–84. 1, 3, 6, 8, 9, 10, 13, 26

[FFT16]   Uriel Feige, Michal Feldman, and Inbal Talgam-Cohen, *Oblivious rounding and the integrality gap*, APPROX-RANDOM, LIPIcs, vol. 60, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016, pp. 8:1–8:23. 2, 6, 9

[GLS81]   M. Grötschel, L. Lovász, and A. Schrijver, *The ellipsoid method and its consequences in combinatorial optimization*, Combinatorica **1** (1981), no. 2, 169–197. MR 625550 5

[GS11]   Venkatesan Guruswami and Ali Kemal Sinop, *Lasserre hierarchy, higher eigenvalues, and approximation schemes for quadratic integer programming with PSD objectives*, Electronic Colloquium on Computational Complexity (ECCC) **18** (2011), 66. 1

[GW01]   Michel X. Goemans and David P. Williamson, *Approximation algorithms for MAX-3-CUT and other problems via complex semidefinite programming*, STOC, ACM, 2001, pp. 443–452. 2

[GZ89]   I. Gilboa and E. Zemel, *Nash and correlated equilibria: Some complexity considerations*, Games Econ. Behav. **1** (1989), 80–93. 1, 2, 4, 14

[HK11]   Elad Hazan and Robert Krauthgamer, *How hard is it to approximate the best Nash equilibrium?*, SIAM Journal on Computing **40** (2011), no. 1, 79–91. 2, 13

[HL18]   Samuel Hopkins and Jerry Li, *Mixture models, robustness and sum-of-squares*, Arxiv (2018). 1

[HNW16]   Aram W. Harrow, Anand V. Natarajan, and Xiaodi Wu, *Tight SoS-degree bounds for approximate Nash equilibria*, Computational Complexity Conference (CCC), 2016. 1, 2, 3, 9, 10

[HSS15]   Samuel B. Hopkins, Jonathan Shi, and David Steurer, *Tensor principal component analysis via sum-of-square proofs*, COLT, JMLR Workshop and Conference Proceedings, vol. 40, JMLR.org, 2015, pp. 956–1006. 1

[KKM18]   Adam Klivans, Pravesh K. Kothari, and Raghu Meka, *Efficient algorithms for outlier-robust regression*, Conference on Learning Theory, COLT (2018). 1

[KMOW17] Pravesh K. Kothari, Ryuhei Mori, Ryan O'Donnell, and David Witmer, *Sum of squares lower bounds for refuting any CSP*, Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017, 2017, pp. 132–145. 1

[KS18a]   Pravesh K. Kothari and Jacob Steinhardt, *Better agnostic clustering via relaxed tensor norms*, Arxiv (2018). 1

[KS18b]   Pravesh K. Kothari and David Steurer, *Outlier-robust moment estimation via sum-of-squares*, Arxiv (2018). 1

[KT10]   Ravi Kannan and Thorsten Theobald, *Games of fixed rank: A hierarchy of bimatrix games*, Economic Theory **42** (2010), no. 1, 157–174, Preliminary version appeared in SODA'07, and available at arXiv:cs/0511021 since 2005. 1, 3

[Las01]   Jean B. Lasserre, *Global optimization with polynomials and the problem of moments*, SIAM Journal on Optimization **11** (2001), no. 3, 796–817. 5

[Li13]   Shi Li, *A 1.488 approximation algorithm for the uncapacitated facility location problem*, Inf. Comput. **222** (2013), 45–58. 3

[LMM03]   Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta, *Playing large games using simple strategies*, ACM Conference on Electronic Commerce, 2003, pp. 36–41. 1, 2, 3, 9

[MP91]   N. Megiddo and C. H. Papadimitriou, *On total functions, existence theorems and computational complexity*, Theor. Comput. Sci. **81** (1991), no. 2, 317–324. 1

[MSS16]   Tengyu Ma, Jonathan Shi, and David Steurer, *Polynomial-time tensor decompositions with sum-of-squares*, CoRR **abs/1610.01980** (2016). 1

[NWY00]   Yuri Nesterov, Henry Wolkowicz, and Yinyu Ye, *Semidefinite programming relaxations of nonconvex quadratic optimization*, Handbook of semidefinite programming, Internat. Ser. Oper. Res. Management Sci., vol. 27, Kluwer Acad. Publ., Boston, MA, 2000, pp. 361–419. MR 1778235 5

[OZ12]   Ryan O'Donnell and Yuan Zhou, *Approximability and proof complexity*, Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, PA, 2012, pp. 1537–1556. MR 3202997 1

[Pap94]   C. Papadimitriou, *On the complexity of the parity argument and other inefficient proofs of existence*, Journal of Computer and System Sciences **48(3)** (1994), 498–532. 1

[Par00]   Pablo A Parrilo, *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*, Ph.D. thesis, Citeseer, 2000. 5

[RRS16]   Prasad Raghavendra, Satish Rao, and Tselil Schramm, *Strongly refuting random csps below the spectral threshold*, CoRR **abs/1605.00058** (2016). 1

[RS09]   Prasad Raghavendra and David Steurer, *How to round any CSP*, 2009 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2009), IEEE Computer Soc., Los Alamitos, CA, 2009, pp. 586–594. MR 2648437 3

[Rub16]   Aviad Rubinstein, *Settling the complexity of computing approximate two-player Nash equilibria*, IEEE Annual Symposium on Foundations of Computer Science, 2016, pp. 258–265. 1, 3

[Sho89]   N. Z. Shor, *Nonsmooth optimization and dual bounds*, Nonsmooth optimization and related topics (Erice, 1988), Ettore Majorana Internat. Sci. Ser. Phys. Sci., vol. 43, Plenum, New York, 1989, pp. 403–410. MR 1034072 5

[SvS06]   Rahul Savani and Bernhard von Stengel, *HardâĂŘtoâĂŘsolve bimatrix games*, Econometrica **74** (2006), no. 2, 397–429. 1

[TS07]   Haralampos Tsaknakis and Paul Spirakis, *An optimization approach for approximate Nash equilibria*, Web and Internet Economics (WINE), 2007, pp. 42–56. 1

[Tul09]   Madhur Tulsiani, *CSP gaps and reductions in the lasserre hierarchy*, STOC, ACM, 2009, pp. 303–312. 13, 14, 21

A Reductions Within Sum-of-Squares

Proof of Fact 5.6. Let $\tilde{\mu}(x)$ be a degree d pseudo-distribution satisfing the constraints of $\Psi_1(x)$. Let $\tilde{\nu}(y)$ be a signed measure on $\mathbb{R}^{n_2}$ defined by the associated expectation functional: $\tilde{\mathbb{E}}_{\tilde{\nu}}[q(y)] = \tilde{\mathbb{E}}_{\tilde{\mu}}[q(f_1(x), f_2(x), \ldots, f_{n_2}(x))$ for every polynomial q. Observe that if the degree of q is a, then the degree of $q(y(x))$ for $y(x) = (f_1(x), f_2(x), \ldots, f_{n_2}(x))$, is at.

We claim that $\tilde{\nu}$ is a degree $\lfloor d/t \rfloor$-pseudo-distribution satisfying the constraints of Ψ_2.

First, observe that since $\tilde{\mathbb{E}}_{\tilde{\nu}}$ is a linear operator, $\tilde{\mathbb{E}}_{\tilde{\nu}}$ is a well-defined linear operator. Let us now verify the remaining properties.

First, consider any polynomial of the form $q = \sum_i r_i g_i$ where r_i are arbitrary polynomials and $deg(r_i) + deg(g_i)$ is at most $\lfloor d/t \rfloor$ in y. Then, $q(y(x))$ has degree at most d in x. Further, for every j, we know that $g_j(y(x)) = 0$. Thus, $\tilde{\mathbb{E}}_{\tilde{\nu}(y)} \sum_i r_i g_i(y) = \tilde{\mathbb{E}}_{\tilde{\mu}(x)} \sum_i r_i g_i(y(x)) = 0$.

Next, consider any polynomial $q = \sum_i r_i g_i + \sum_i s_i h_i$ for arbitrary polynomials r_i and sum-of-squares polynomials s_i such that each term appearing in the expansion has syntactic degree of at most $\lfloor d/t \rfloor$.

We have: $\tilde{\mathbb{E}}_{\tilde{\nu}(y)} \sum_i r_i g_i + \sum_i s_i h_i = \tilde{\mathbb{E}}_{\tilde{\mu}(x)} \sum_i r_i(y(x)) g_i(y(x)) + \sum_i s_i(y(x)) h_i(y(x))$.

Now, $g_i(y(x)) = 0$ for every i. Further, $h_i(y(x)) = Q_i(y(x))$ for some sum-of-squares polynomial Q in y. This also means that $Q_i(y(x))$ is a SoS polynomial in x. Finally, since s_i is a SoS polynomial in y, $s_i(y(x))$ is a SoS polynomial in x. Finally, product of SoS polynomials is SoS, thus, $s_i(y(x)) h_i(y(x))$ must also be SoS.

Now, using that $\tilde{\mathbb{E}}_{\tilde{\mu}(x)}$ is a valid pseudo-expectation of degree d completes the proof. $\qquad\square$